

SCHOOL OF OPERATIONS RESEARCH  
AND INFORMATION ENGINEERING  
COLLEGE OF ENGINEERING  
WEILL MEDICAL COLLEGE  
CORNELL UNIVERSITY  
ITHACA, NEW YORK 14853-3801

TECHNICAL REPORT NO. 1471

March 2009

**Dynamic Modeling of Workforce Requirements for Mass Prophylaxis Under Highly  
Uncertain Conditions: Cornell Dynamic POD Simulator**

by

Michelle Castorena  
Caitlin Hawkins  
Nathaniel Hupert  
Kathleen King  
John Muckstadt  
Cindie Wu  
Wei Xiong

## Table of Contents

<b>ABSTRACT .....</b>	<b>3</b>
<b>I. INTRODUCTION.....</b>	<b>4</b>
<b>II. CURRENT RESEARCH .....</b>	<b>5</b>
<b>III. METHODS .....</b>	<b>6</b>
MODEL DESCRIPTION .....	7
MODEL PARAMETERS .....	7
EXPERIMENTAL DESIGN .....	9
<b>IV. RESULTS AND DISCUSSION.....</b>	<b>10</b>
1. HOW IMPORTANT IS HAVING STAFFING MATCH A (NON-STATIONARY) ARRIVAL RATE PATTERN? .....	10
2. WHAT HAPPENS TO POD PERFORMANCE WHEN STAFFING LEVEL ADJUSTMENTS ARE MADE AFTER SOME DELAY IN RESPONSE TO OBSERVED ARRIVAL PATTERNS? .....	15
3. HOW ARE OPTIMAL STAFFING LEVELS AFFECTED BY THE NUMBER OF PODS IN A POD NETWORK?.....	23
4. WHAT ARE THE CONSEQUENCES OF USING A COMMAND AND CONTROL SYSTEM TO MONITOR AND CONTROL THE POD NETWORK? .....	24
<b>V. CONCLUSIONS .....</b>	<b>25</b>
<b>REFERENCES .....</b>	<b>26</b>
<b>APPENDIX B. USER’S MANUAL.....</b>	<b>27</b>
1. INTRODUCTION .....	27
2. GLOSSARY OF IMPORTANT TERMS.....	28
3. WORKING WITH D-PODS .....	29
3.1 <i>Getting Started</i> .....	29
3.2 <i>Starting the Model</i> .....	30
3.3 <i>D-PODS Menu</i> .....	31
3.4 <i>Step 1: Constructing the Model</i> .....	33
3.5 <i>Step 2: Input Arrival Rate Information</i> .....	35
3.6 <i>Step 3: Service Time Parameters</i> .....	37
3.7 <i>Step 4: Staffing Requirements</i> .....	39
3.8 <i>Step 5: Simulation Parameters</i> .....	41
3.9 <i>Step 6: Input Case Name</i> .....	41
3.10 <i>Step 7: Run the Simulation</i> .....	42
3.11 <i>Step 8: View the Results</i> .....	43
4. NAVIGATING THE ACCESS DATABASE .....	48
5. DESIGNING POD STAFFING POLICY EXPERIMENTS.....	49
6. REFERENCES .....	50
<b>APPENDIX C. FUTURE EVENT LISTS .....</b>	<b>51</b>

## Abstract

The public health response to a bioterrorist attack or other large-scale health emergency may include mass prophylaxis using multiple Points of Dispensing (PODs) to deliver countermeasures rapidly to affected populations. Although computer models exist to determine "optimal" staffing levels at PODs under certain steady-state conditions, no quantitative studies address the requirements of the POD workforce management systems needed to enable efficient, population-wide coverage in the face of a dynamic and uncertain operational environment.

Our goal was to investigate quantitatively the impact of dynamic and uncertain patient arrival patterns on workforce requirements and on overall POD effectiveness and efficiency over the duration of a mass prophylaxis campaign.

To investigate the dynamic behavior of POD systems, many extensive simulation experiments were conducted using a Monte Carlo simulation model, called the Dynamic POD Simulator (D-PODS). Using this simulation environment, we designed POD station layouts, capacities, staffing patterns, patient types, and work flows and then observed the consequences of operating PODs based on these plans under various patient arrival scenarios. The D-PODS user interface is an Excel worksheet and the model is implemented in Visual Basic.

Using several illustrative POD experiments, we demonstrate that uncertain patient arrival patterns require higher staffing levels than might be expected if a stationary environment is assumed. These experiments further show that PODs may develop severe bottlenecks unless staffing levels vary over time to meet changing patient arrival patterns. Because of the unpredictability of the operating environment, efficient POD networks require command and control systems capable of dynamically adjusting intra- and inter-POD staffing levels to meet demand. Furthermore, we show that fewer large PODs require a smaller total staff than many small PODs require to serve the same number of patients.

We conclude that modeling environments that capture the effects of fundamental uncertainties in public health disasters are essential for the realistic evaluation of response mechanisms and policies. D-PODS quantifies POD operational efficiency under more realistic conditions than have been modeled previously. Our experiments demonstrate the critical role of variation and uncertainty in POD arrival patterns in establishing effective POD staffing plans. These experiments also highlight the need for command and control systems to be created to manage emergency response successfully.

## I. Introduction

When a bioterrorist attack or other large-scale public health emergency occurs, planned health system response procedures are put into operation.<sup>2</sup> One possible response is the opening of dispensing clinics at which a portion of or the entire population of a given geographical region will receive medications or vaccinations within a very short period of time, perhaps within 48 hours of the time that officials decide to conduct the mass prophylaxis campaign.<sup>3</sup> These clinics are called Points-of-Dispensing, or simply PODs, and cities such as New York have plans to activate hundreds of PODs for some emergencies. The efficiency and effectiveness of the POD system depends crucially on many factors.<sup>4</sup>

Public health planners must determine how many PODs to deploy and where to locate them. Planners must also carefully design each individual POD. In many POD layouts, the tasks associated with dispensing medication (collecting patient information, prescribing a dosage or medication, packaging drugs, etc.) are divided among several service stations. Patients move from station to station in one of a few possible paths, depending on their specific conditions and circumstances. The effectiveness of this type of POD plan depends on the accuracy of assumptions made about the lengths of time required to process different patient types (families of different sizes, elderly people, disabled individuals, etc.), the percentages of various patient types that follow each possible route through the POD, and the pattern of arrivals throughout the duration of POD operations. Staffing levels for each service station over time will be determined based on these assumptions. When establishing these staffing levels, planners must recognize that their assumptions may not be correct. Real events are usually very difficult to predict with accuracy.

Our goal, in general, was to explore the consequences of POD staffing policies. To do this, we created a flexible simulation environment called the Dynamic POD Simulator (D-PODS) that enabled us to understand the dynamics of patient flow within a POD in a probabilistic way. This, in turn, allowed us to consider many potential outcomes from a mass prophylaxis campaign. In the following sections we will describe other existing POD models and the simulation model that we implemented before proceeding to discuss a number of questions related to POD staffing policies.

First, we show that staffing levels must vary over time to meet changing patient arrival patterns. Second, we show that, faced with uncertainty in arrival patterns, patient waiting times quickly become unreasonable unless staffing levels are either extremely flexible or significantly larger

than might be expected. Third, we demonstrate that the total number of staff needed to run a network of PODs depends critically on the number of PODs in the system. Specifically, we show that to serve the same number of patients, fewer larger PODs require a smaller total staff than many smaller PODs. Finally, we show that for a prophylaxis campaign to be successful, a POD network requires a system capable of dynamically managing its operations. To mitigate the effects of uncertainty, planners must create a command and control environment in which demand for services at each POD can be assessed and responded to frequently. It also means that staff may have to be moved between PODs or patients transported to PODs with “lighter” demand (or both) at various points in time for a mass prophylaxis campaign to be completed effectively. While we will primarily discuss how D-PODS can be used to set POD staffing policy, the observations we make can also be applied to the allocation of other resources, such as medical supplies and equipment.

## II. Current Research

A number of research groups have developed approaches for designing and evaluating the performance of alternative POD systems. Three of the most detailed systems have been developed at Weill Cornell Medical College, Georgia Institute of Technology, and the University of Maryland. These research efforts provided us with insights into the design and implementation of D-PODS.

At the Weill Cornell Medical College, Hupert et al. developed BERM (Bioterrorism and Epidemic Outbreak Response Model, available at <http://www.simfluenza.org>), a tool that is widely used to calculate POD staffing levels.<sup>5</sup> BERM’s output is based on the assumption that patients arrive at the POD according to a stationary Poisson process. Using a Jackson queuing network model and an optimization model, BERM calculates the minimum staffing levels needed at each station to achieve a desired level of performance. Although these results are beneficial to planners, it is highly unlikely that patients will arrival at PODs in a stationary manner. In D-PODS we allow users the option of setting non-stationary patient arrival rates.

Eva Lee et al., at the Georgia Institute of Technology, developed RealOpt, which is a tool that can be used to evaluate staffing levels and develop a customized layout for a POD.<sup>6</sup> The tool uses a heuristic-algorithmic design and works in real-time as a decision support system during an actual dispensing campaign. RealOpt aids in strategic and operational planning before, during, and after an emergency mass-dispensing event, but primarily focuses on the ability to reconfigure staffing levels during operations. The assumption that a POD can be easily

redesigned and staff reallocated at any given moment may be unrealistic. In theory, reconfiguration of a POD may increase its efficiency. However, if staff members are to be moved among POD stations, they must be trained in advance for the many different jobs they might be expected to perform; otherwise the potential gain in efficiency could be lost in training workers in new tasks during the emergency.

At the University of Maryland, Aaby and others have used the results of a small time study and a simulation model implemented in Arena to develop the Clinic Planning Model Generator (CPMG), a queuing model-based capacity-planning tool which estimates POD performance based on user-specified inputs.<sup>7</sup> Aaby et al. used the results of the time study to validate their simulation model, and then evaluated different clinic designs and operational policies using this model. The model allows patients to arrive either individually or in batches (as if people came by buses or children came with school classes), and Aaby et al. studied the effect of changing the distribution of the inter-arrival times of these buses. CPMG also allows for the creation of stations capable of serving an entire batch of patients at the same time, and the researchers considered a variety of batch sizes. CPMC was created to quickly provide accurate suggestions on appropriate staffing levels and to analyze the performance of the clinic based on those suggestions. One powerful aspect of CPMG is its ability to create a customized clinic layout. However, like the original BERM, the model assumes stationary arrival rates when it calculates the staff required for each station, patient throughput, patient time spent in the system, and the effect of changing the number of stations. Thus, while CPMG is helpful in evaluating different clinic designs and suggesting layout and operational guidelines, its output is potentially unreliable because stationary patient arrival rates are unlikely in a real emergency.

Despite their drawbacks, all of the tools discussed above have proven to be useful in planning and operating mock prophylaxis exercises. They have provided POD planners with the means to calculate required resources and to determine results of certain staffing level decisions. While all these tools help us to gain insight into the POD operations, they do not sufficiently represent the incredible degree of uncertainty that would be present in an emergency situation. The model we next describe was created to assess the effects of this uncertainty.

### **III. Methods**

We developed a simulation model of a POD that we used to experiment with a variety of staffing policies and patient arrival scenarios. In particular, we considered scenarios of a POD responding

to a large-scale airborne anthrax attack. Below we describe the model and its assumptions, the simulation parameters used, and, most importantly, the policy-related questions that we explored.

### ***Model Description***

In our POD model, there is a set of stations at which incoming patients are served. Patients of various types (single adult, family, limited mobility, etc.) arrive to the POD according to nonhomogeneous Poisson processes. They move from station to station according to transition probabilities. The number of staff at each station may vary over the course of simulated time horizon. At each station, patients wait in a queue until a staff member is free, at which point the patient receives service. Once the patient has received service, he immediately joins the queue for the next station or departs from the POD; we assume that patient travel time within a POD is negligible. The entire POD can hold only a limited number of patients, so if more people arrive than can fit in the POD, we assume that they form a queue outside. We assume that no patients abandon their queues, either inside or outside of the POD, no matter how long they have waited. We further assume that once a patient enters the POD, a patient will receive service.

Patient service times at each station within the POD are distributed according to triangular probability distributions specific to the station. Hupert et al. used an exponential distribution for modeling service times.<sup>4</sup> However, exponential distributions have large tails, which could allow extremely large service times, when, in practice, an upper bound on the service time is more realistic. In the anthrax dispensing exercise performed by Lee et al., triangular probability distributions were observed to represent service times accurately.<sup>6</sup> Another assumption of the model is that each server processes one patient at a time. Batching of patient arrivals is left for a later study. Furthermore, D-PODS does not yet account for a staff learning curve or fatigue. The staffing levels include only the core staff (servers at each station), and calculations for support staff have not yet been implemented.

We implemented the model described using a highly flexible, user-friendly tool, called the Cornell Dynamic POD Simulator (D-PODS). For more information regarding the tool, see Appendix B.

### ***Model Parameters***

To illustrate how the simulation model can be used to assess the effectiveness of POD system policies, we have chosen parameters to reflect a network that is set up to respond to an airborne

anthrax attack. Studies have shown that antibiotic dispensing to prevent symptomatic inhalational anthrax is most effective if administered within 3 days of exposure.<sup>9-13</sup> Since it will take some time for the attack to be detected and a decision made to begin a prophylaxis campaign, the CDC goal, as stated in the Cities Readiness Initiative, is for antibiotics to be dispensed to the entire affected population within 48 hours of the decision to do so.<sup>15</sup> Thus, we set the duration of the prophylaxis campaign to 48 hours.

We used a standard POD layout, consisting of 4 stations: Greeting, Triage, Medical Evaluation, and Drug Dispensing. The layout diagram and routing probabilities are shown in Figure 1.

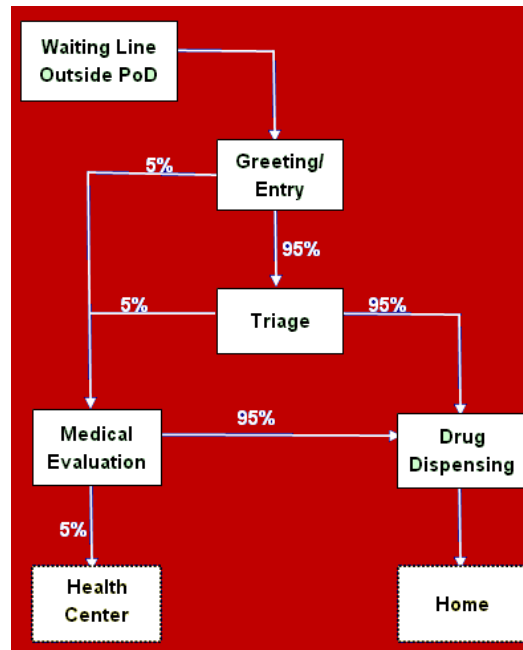


Figure 1. POD patient flow diagram.

For simplicity, we assume that the POD is large enough to accommodate all arriving patients, although D-PODS does permit users to limit the POD's capacity. We assume that the POD closes only for two hours each day, during which restocking and cleaning occur. During the 22 hours when patients are arriving, the POD is expected to serve 11,000 patients, or about 500 people per hour if on average people arrive at a constant rate throughout the day. To further simplify our discussion, we assume that all arriving patients are of the same "type," which means that the service time distributions are the same for all patients. These service times and other parameters are summarized in Table 1.

Duration of Campaign	2	days
Hours of operation per day	22	Hours
Probability that a patient at the Greeting station will	0.05	(n/a)



require Medical Evaluation		
Probability that a patient at the Triage station will require Medical Evaluation	0.05	(n/a)
Probability that a patient will be sent to a Health Center after Medical Evaluation	0.05	(n/a)
Service time distribution for Greeting/Entry station	Triangular(0.25,0.50,1)	minutes
Service time distribution for Triage station	Triangular (1,2,3)	minutes
Service time distribution for Medical Evaluation station	Triangular(2.5, 5, 10)	minutes
Service time distribution for Drug Dispensing station	Triangular (0.5, 1, 2)	minutes
Number of arrival types	1	type
Service time increase factor	0%	(n/a)

**Table 1. Simulation parameters.**

## ***Experimental Design***

Our experiments were designed to answer four essential questions regarding POD effectiveness under highly uncertain conditions:

1. How important is having staffing match a (non-stationary) arrival rate pattern?
2. What happens to POD performance when staffing level adjustments are made after some delay in response to observed arrival patterns?
3. How are optimal staffing levels affected by the number of PODs in a POD network?
4. What are the consequences of using a command and control system to monitor and control the POD network?

In answering these questions, we simulated a number of different arrival scenarios and staffing plans. For the first question, we compared what happens when staffing levels are constant throughout the day versus when they change frequently in response to various patient arrival patterns. In response to the second question, we considered more complex staffing strategies in which we adjust staffing periodically throughout the day in response to observed patient arrivals. For the third question, we ran simulations of large and small PODs and calculated the number of staff required given the number of operating clinics. The simulation results obtained in response to the first three questions also provide sufficient evidence to allow us to answer the fourth question.

Note that in each of the figures shown, we display only 10 simulation repetitions. We did this to provide clear, easily-readable figures, since our purpose here is to show general principles rather than to determine strict numerical results. While larger numbers of replications give D-PODS results greater statistical significance, the conclusions remain unaltered.

## IV. Results and Discussion

### ***1. How important is having staffing match a (non-stationary) arrival rate pattern?***

As mentioned earlier, POD planning tools most often are based on the assumption that patient arrivals occur at a constant rate during all POD operating hours. However, in the event of an actual emergency, patient arrival rates will likely vary throughout the day. If people go to work during the emergency, then we might expect to see peaks in patient arrival rates at the beginning and end of the work day. We might also see spikes in arrivals due to changes in the availability of public transit, public health department publicity encouraging people to arrive at different times, or at the beginning of the school day, if children are brought to the PODs via their schools. Peaks and valleys in the arrival process may also arise randomly throughout the day. Thus, we want to show that POD staffing policies should be designed so that PODs will be able to respond to arrival rates that vary throughout a prophylaxis campaign.

To accomplish this objective we will examine what happens when staffing levels remain constant throughout each day, versus what happens when staffing is adjusted for actual patient arrival patterns. We demonstrate that, if we can predict patient demand correctly, then patients will spend far less time in the POD, staff will neither be under- nor over-utilized, and queue lengths at all of the POD stations will be far more reasonable than when staffing levels are held constant throughout each day. To show this, we considered three potential daily patient arrival patterns; Figure 2 shows the expected hourly arrival rates for each pattern. For each arrival pattern, we used D-PODS to compare POD operations with constant staffing levels versus dynamic staffing levels that were allowed to change every two hours throughout the day. For the dynamic staffing plans, we determined staffing levels during each two hour period by calculating the number of staff needed to limit average patient waiting times at each station to five minutes or less. The constant staffing levels were estimated by calculating the number of staff required to limit average patient waiting times at each station to five minutes or less, assuming that the hourly patient arrival rate is constant throughout the day. In both cases, the total expected number of daily arrivals was 11,000, as mentioned in Table 1.

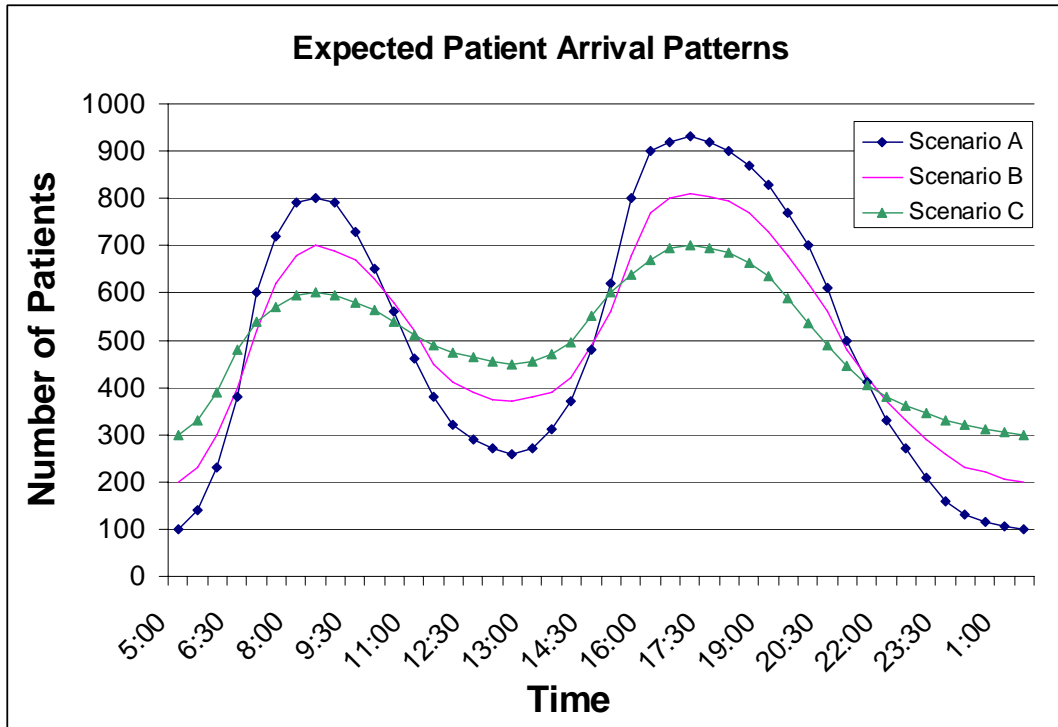


Figure 2: Expected Patient Arrival Scenarios.

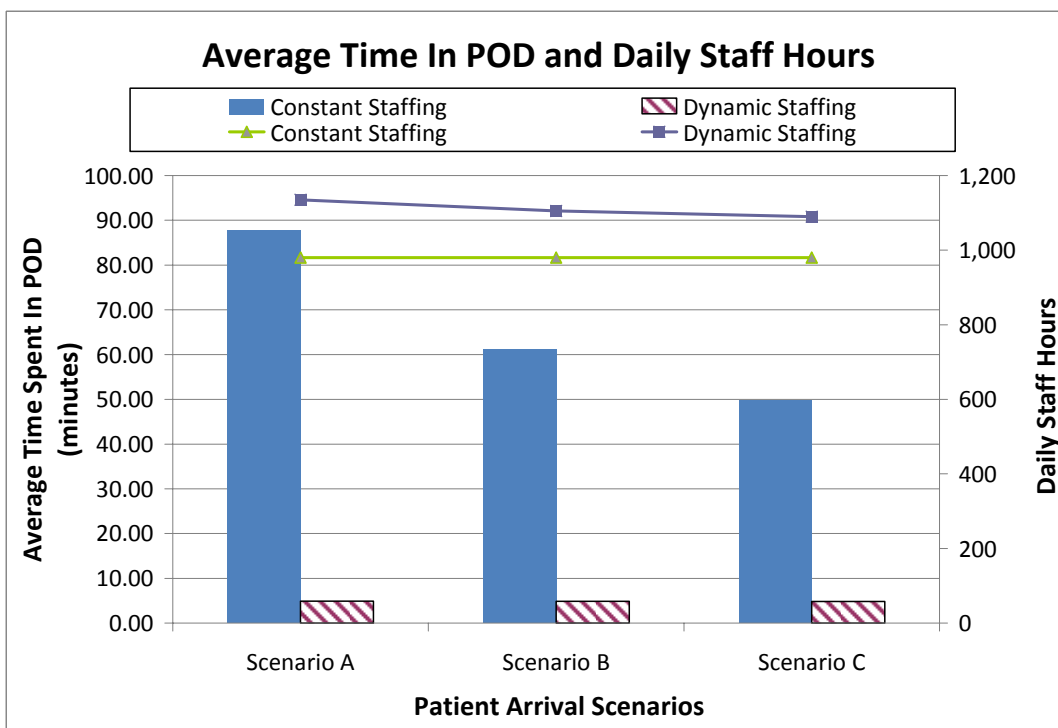
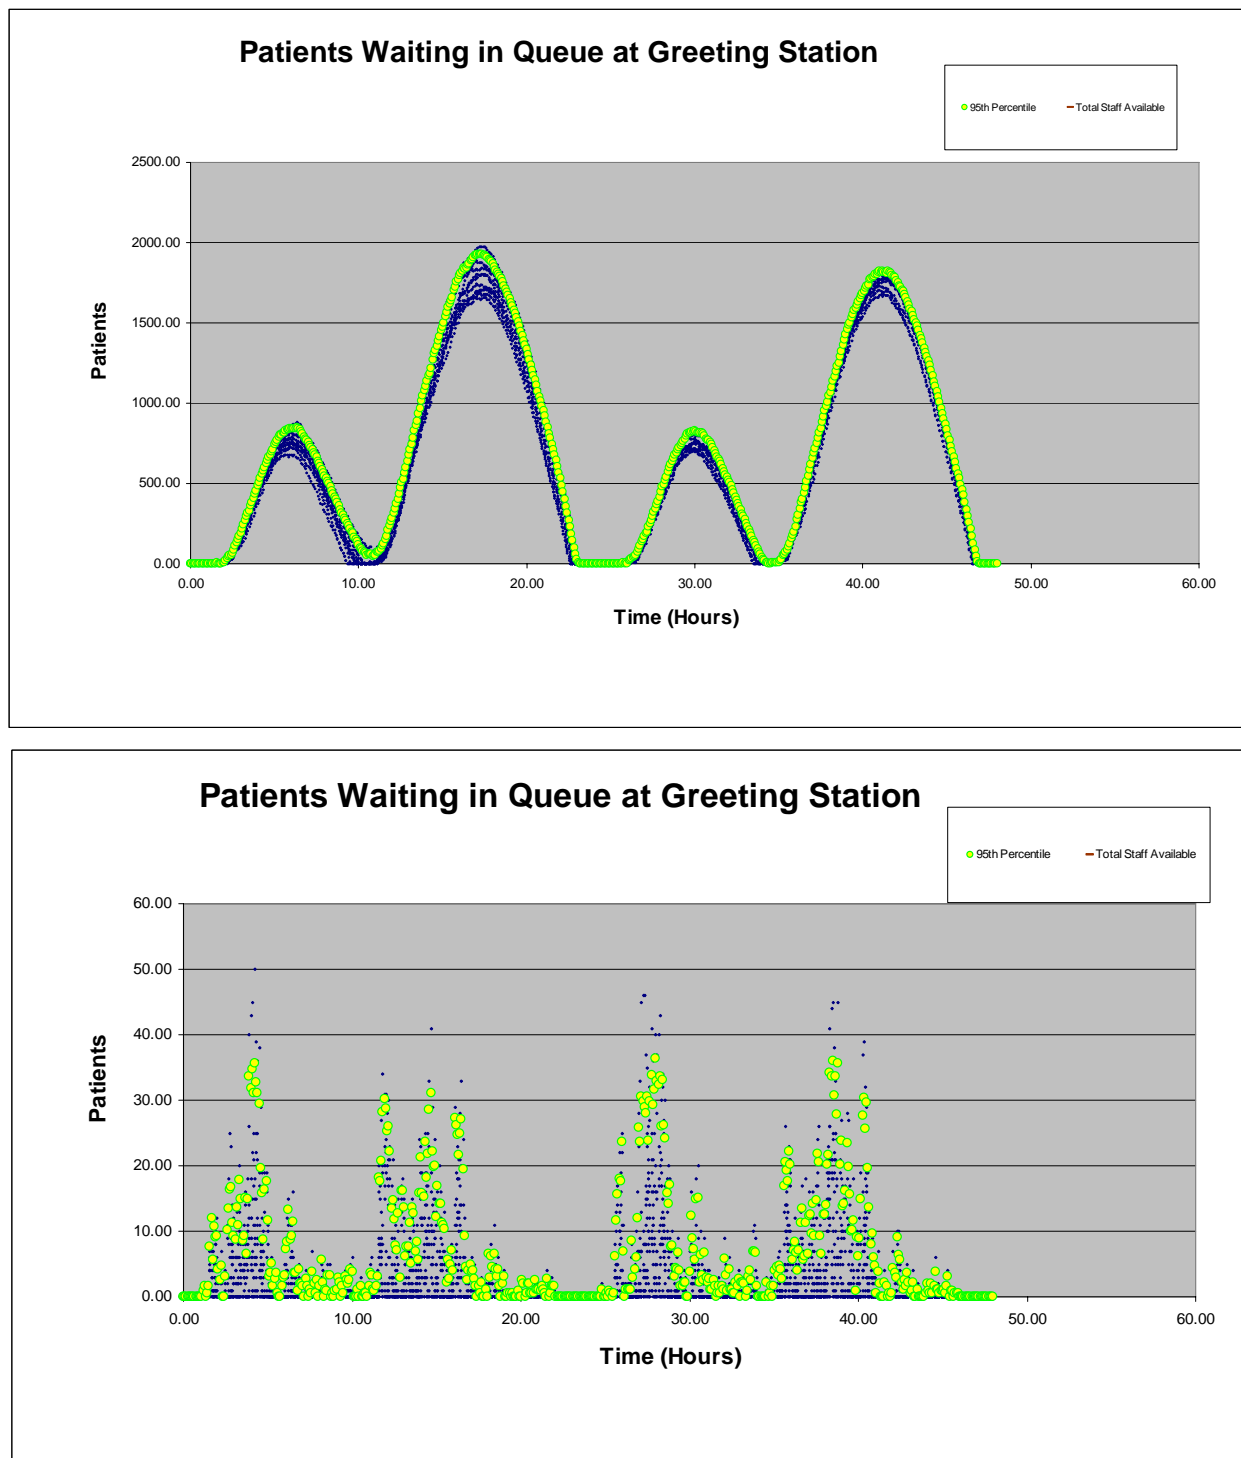


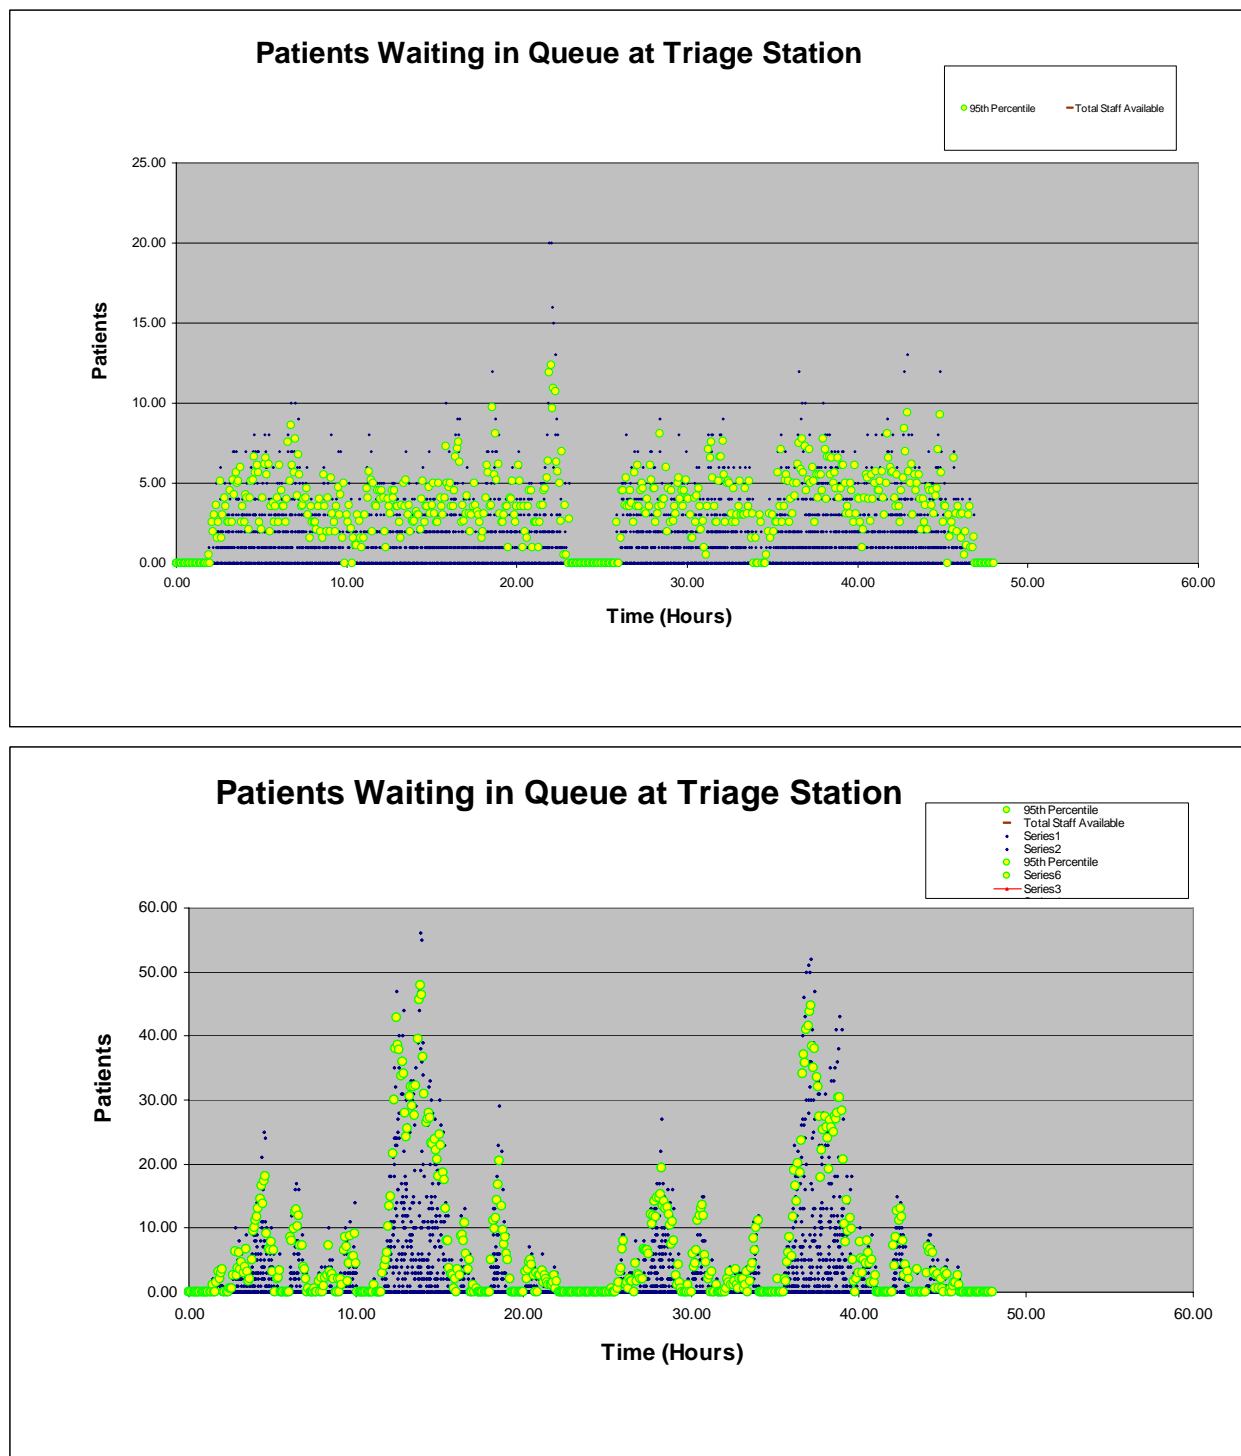
Figure 3: Average Patient Time Spent in the POD and Staff-Hours Each Day, given Constant or Dynamic Staffing Plans.

In Figure 3, we see that the average patient waiting time vastly decreases when the staffing plan is adjusted to correspond to the dynamic patient arrival pattern. In all three scenarios, the average patient time spent in the POD was only about five minutes for the dynamic staffing plans, while the average time for the constant staffing plans ranged from 50 minutes to almost 90 minutes. It must also be noted that the dynamic staffing plans required 10% to 14% more staff-hours than the constant staffing plan. In general, given a total daily throughput and a limit on the maximum patient waiting time at each station, if patients arrive at a constant rate during each day then staffing requirements will be lower than those needed when patient arrival patterns vary over time.

The averages shown in Figure 3 cannot fully describe the dynamics of POD operations, of course. Consider Figures 4 and 5, which show the actual queue lengths that arose in the simulation replications for the Greeting and Triage stations, for Scenario A. The top graph in each figure shows the results from using the constant staffing plan; notice the different scales between the two graphs. The queue at the Greeting station becomes unstable for every peak of the arrival process, growing larger than 1200 people for every simulation replication. There are only 5 servers at the Greeting station, so one could easily imagine the frustration of arriving patients upon seeing such a huge, slow-moving line. In the real world, it is possible that many individuals would give up and go home before receiving treatment, which could be disastrous in an anthrax attack had occurred.



**Figure 4: The length of the queue at the Greeting Station for Scenario A for Constant Staffing (top) and Dynamic Staffing (bottom).** The small blue dots show the queue lengths from the 10 simulation replications for every five minute interval; the yellow-green circles show the 95<sup>th</sup> percentile queue length for each five minute interval.



**Figure 5: The length of the queue at the Triage Station for Scenario A for Constant Staffing (top) and Dynamic Staffing (bottom).** The small blue dots show the queue lengths from the 10 simulation replications for every five minute interval; the yellow-green circles show the 95<sup>th</sup> percentile queue length for each five minute interval.

Under the dynamic staffing plan, which varies corresponding to patient arrival patterns, the Greeting station queue almost always remains smaller than 10 or 20 people, which is clearly a much more desirable situation. However, notice in Figure 5 that the queue at the Triage station is often smaller under the constant staffing plan than the dynamic staffing plan. This happens because the bottleneck at the Greeting station under the constant staffing plan slows the flow of patients to a trickle at the Triage station; thus, the Triage staff are never overwhelmed. Under the dynamic staffing plan, however, queues and delays are balanced between the stations. One possible option for a POD with a bottleneck such as the one that occurs here under the constant staffing plan would be to move staff from the Triage station to help out at the Greeting station. Of course, if there are simply not enough staff in the entire POD, moving staff will have the effect of shifting the bottleneck from Greeting to Triage. Only when the system has excess capacity can dynamic shifting of staff have a long-term positive impact on patient waiting times.

The results from Scenarios B and C are similar; using a dynamic staffing plan prevents large queues at the Greeting station, and a bottleneck occurs at the Greeting station under the constant staffing plan, which limits the queue length at the Triage station. However, the Greeting station queue lengths under the constant staffing scenario are shorter in Scenario B than in Scenario A, and they are shorter still in Scenario C. This is not surprising since the patient arrival rates in Scenarios B and C are less variable over time than the rates in Scenario A. In short, when patient arrival rates are closer to being constant, using a constant staffing plan does not increase waiting times as significantly.

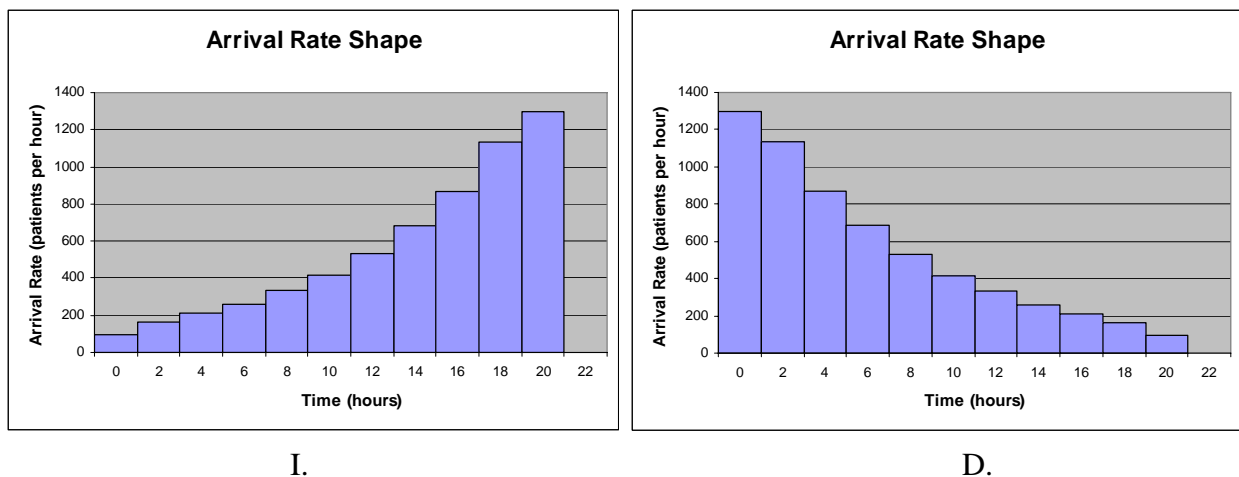
These results show that staffing PODs based on actual time-varying patient arrival patterns is extremely important. However, staffing in this manner requires that planners estimate expected patient arrival patterns exactly, and such perfect prediction is highly unlikely. It is similarly unlikely that staff will be able to arrive or leave the POD instantaneously when patient arrival rates change. Thus, in the next section we consider what happens when staffing changes are made periodically throughout the day when there are delays in staffing adjustments and predictions of patient arrivals are inaccurate.

## ***2. What happens to POD performance when staffing level adjustments are made after some delay in response to observed arrival patterns?***

We now consider staffing policies in which a number of staff are on-call, but they are brought in or sent home throughout the day, according to patient demands. Staffing levels, then, are determined by forecasted demand levels calculated dynamically throughout the day. We will

suppose that staffing levels may change as often as every two hours. We will consider the several staffing plans.

In Plan 1, we suppose that we are able to perfectly forecast expected patient arrival rates throughout each day and we can change staffing levels at each station every two hours. Then, as in the dynamic staffing plan from the previous question, we use a queuing formula from Buzacott<sup>16</sup> to conservatively estimate workforce requirements at each station to ensure that the average patient waiting time at each station will not likely exceed five minutes. In Plan 2, we suppose that one half hour before each staffing change, we could observe the true average patient arrival rate to the POD at that time. We assume that patients will, on average, continue to arrive at this rate for the next worker interval. We use the same queuing formula mentioned above to determine the staffing levels that would be required to limit the average patient waiting time to five minutes at each station. Plan 3 is similar to Plan 2, except that arrival rates must be observed and staffing decisions made 60 minutes before each worker interval. This means that we assume that, during the worker interval, patients will arrive at the average rate that is seen one hour before the worker interval begins. In Plan 4, we suppose that the planner is able to correctly forecast the total number of arrivals for the next two hours, and she sets staffing levels to limit expected patient waiting times to five minutes at each station, assuming that the expected patient arrival rate is constant for the next two hours. For a more detailed explanation on how these staffing plans are constructed, see Appendix A.



**Figure 6. Expected Arrival Rate Patterns that are increasing (I) or decreasing (D).**

For simple patient arrival patterns like those shown in Figure 6, it is easy to predict the results of using a staffing policy in which the future patient arrivals are forecasted solely based on the



current patient arrival rate, like Plans 2 and 3. Let us assume that staffing levels are set correctly for the first hour of the day. If patients arrive according to pattern (I), then the planned staffing levels will always be too low. The longer it takes to adjust staffing levels, the longer the queues and patient waiting times will become at the POD. If patients arrive according to pattern (D), then the planned staffing levels will always be too high, and the longer it takes to adjust staffing levels, the more under-utilized the staff will become. This consequence may sound insignificant compared to huge queues and excessive patient waiting times, but if you only have a limited number of total available staff-hours, this type of policy could mean that the POD will be understaffed when demand is high.

We now consider patient arrival Scenario A described in the previous section. In Figure 7 we see the average patient waiting times for each of the plans. As we might have expected, Plan 3 gives the longest patient waiting times, but notice that these are still significantly better than the wait times observed for the constant staffing plan discussed in the previous section (see Figure 3). Dynamically adjusting staff, even using a forecasting method like those in Plans 2 and 3, decreases waiting times on average by almost 30 percent.

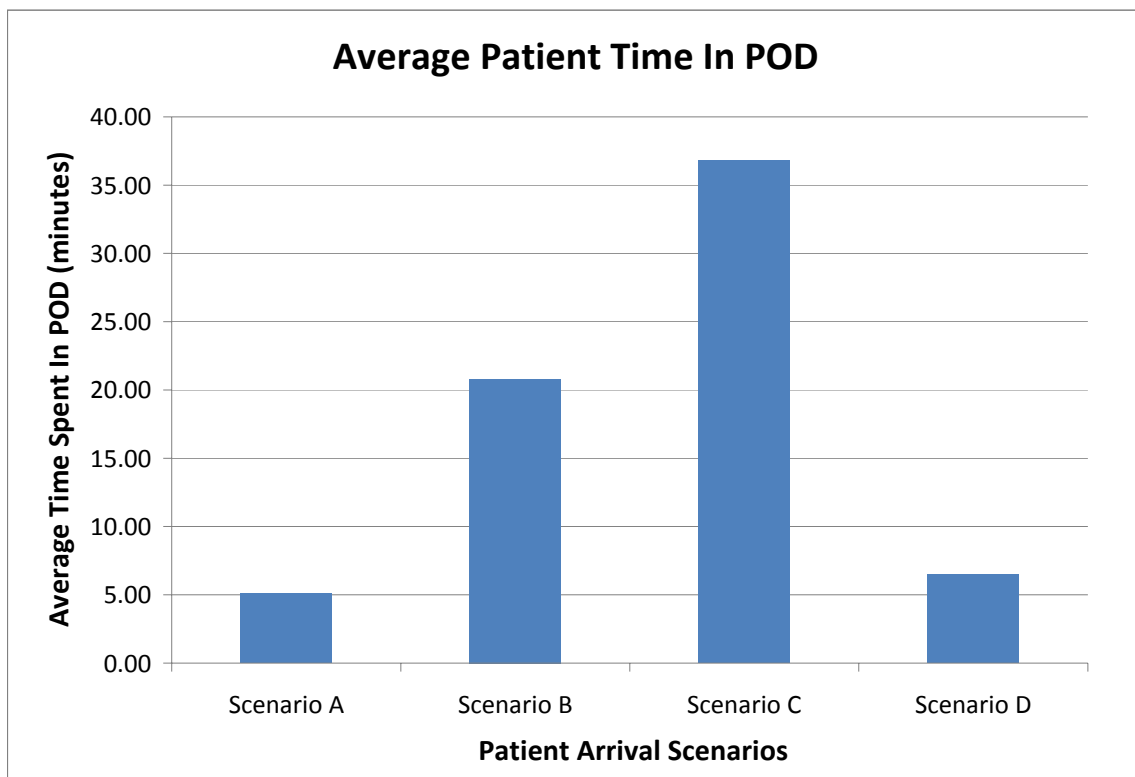


Figure 7. Average Patient Time In POD for Arrival Scenario A.

Figure 7 may seem to imply that Plans 1 and 4 are almost equally good. This would be excellent news, since the staffing levels used in Plan 4 are far less closely tuned to patient arrivals than the staffing levels used in Plan 1. However, Figures 8 and 9 show the dynamics of these two simulations, and we see that the queues at both the Greeting and Triage grow to significant lengths under Staffing Plan 4. At the Greeting Station, the queue grew longer than 100 people in many of the replications under Plan 4, but under Plan 1 the queue seldom crept above 30. At the Triage Station we see that the queue became very unstable and grew longer than 200 in one or two replications under Plan 4, but generally remained stable under Plan 1.

We show these results only for patient arrival Scenario A, which is the most “extreme” of the three patient arrival patterns considered. As might be expected, the performance of Plans 2, 3, and 4 is better under the more “moderate” patient arrival patterns, Scenarios B and C. In fact, notice that all four plans would be identical if patient arrival rates were constant. If patient arrival rates changed exactly when staffing shifts changed and stayed constant during each worker interval, then Plans 1 and 4 would provide identical staffing levels. In general, Plans 1 and 4 will be similar, but Plan 1 will provide slightly better POD performance and Plan 4 will provide slightly higher staff utilization and lower total staff-hours. This occurs because Plan 1 always staffs for the maximum expected patient arrival rate during each worker interval while Plan 4 staffs for the average expected patient arrival rate during each worker interval.

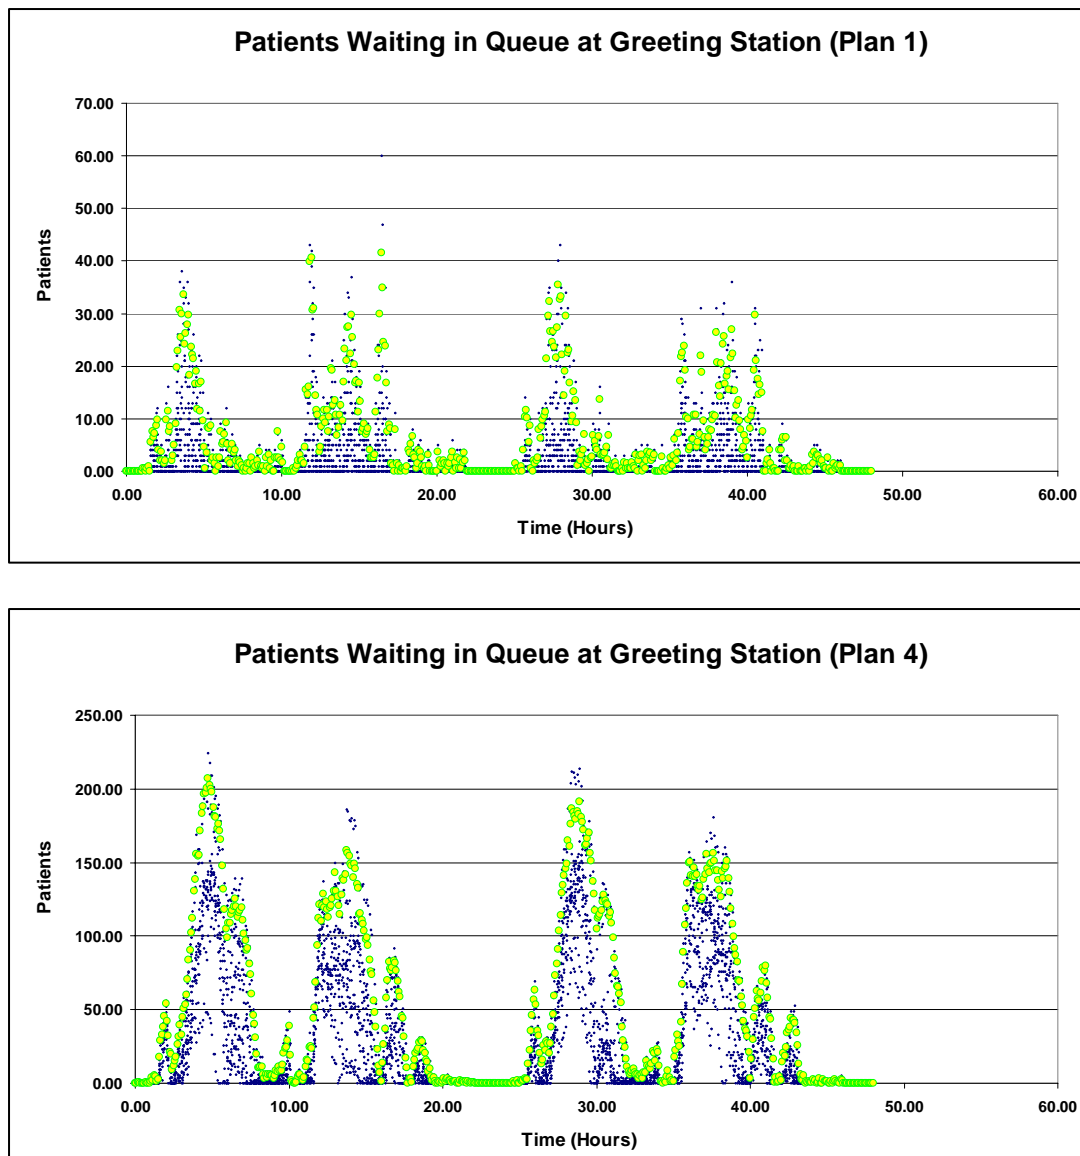
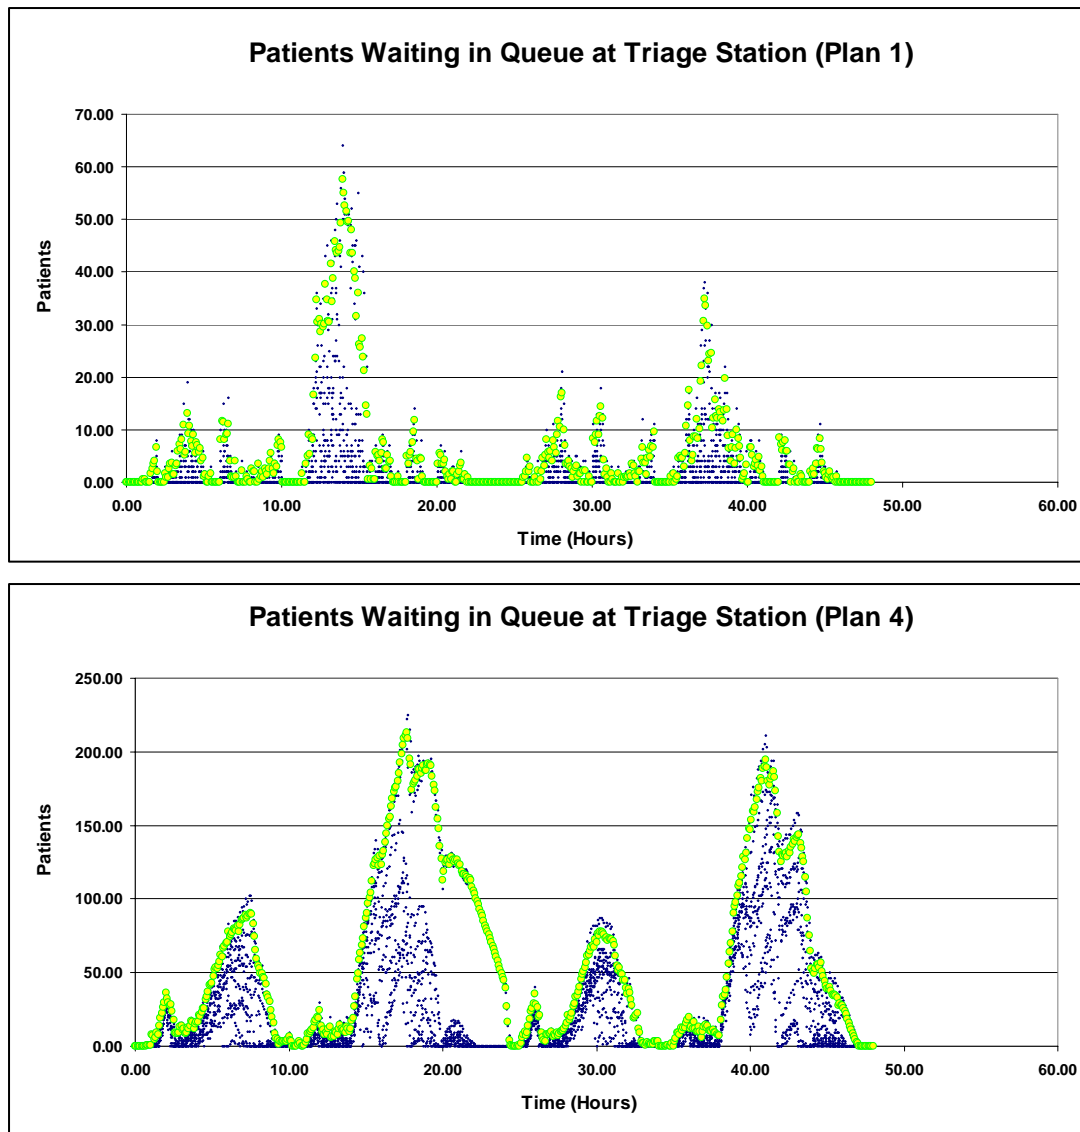


Figure 8: Queue Lengths at the Greeting Station for Plans 1 and 4 With Patient Arrival Scenario A.



**Figure 9: Queue Lengths at the Triage Station for Plans 1 and 4 With Patient Arrival Scenario A.**

It is also important to consider staffing plans that assume that patient arrival rates will remain approximately constant over short time intervals. Plans 2 and 3 are both of this type, and in Figure 7 they appear to be acceptable, since 20 or 35 minutes is not an unreasonable amount of time to require patients to wait for service. The more detailed picture seen in Figure 10 is significantly bleaker. In Figure 10 we see the queue lengths at three of the POD stations using Plan 3 for the scenario with the lowest variation, namely Scenario C. We do not show results for Scenario B, since they are simply intermediate compared to the results from Scenarios A and C, so to limit repetition we omit the plots from the Scenario B simulations here.

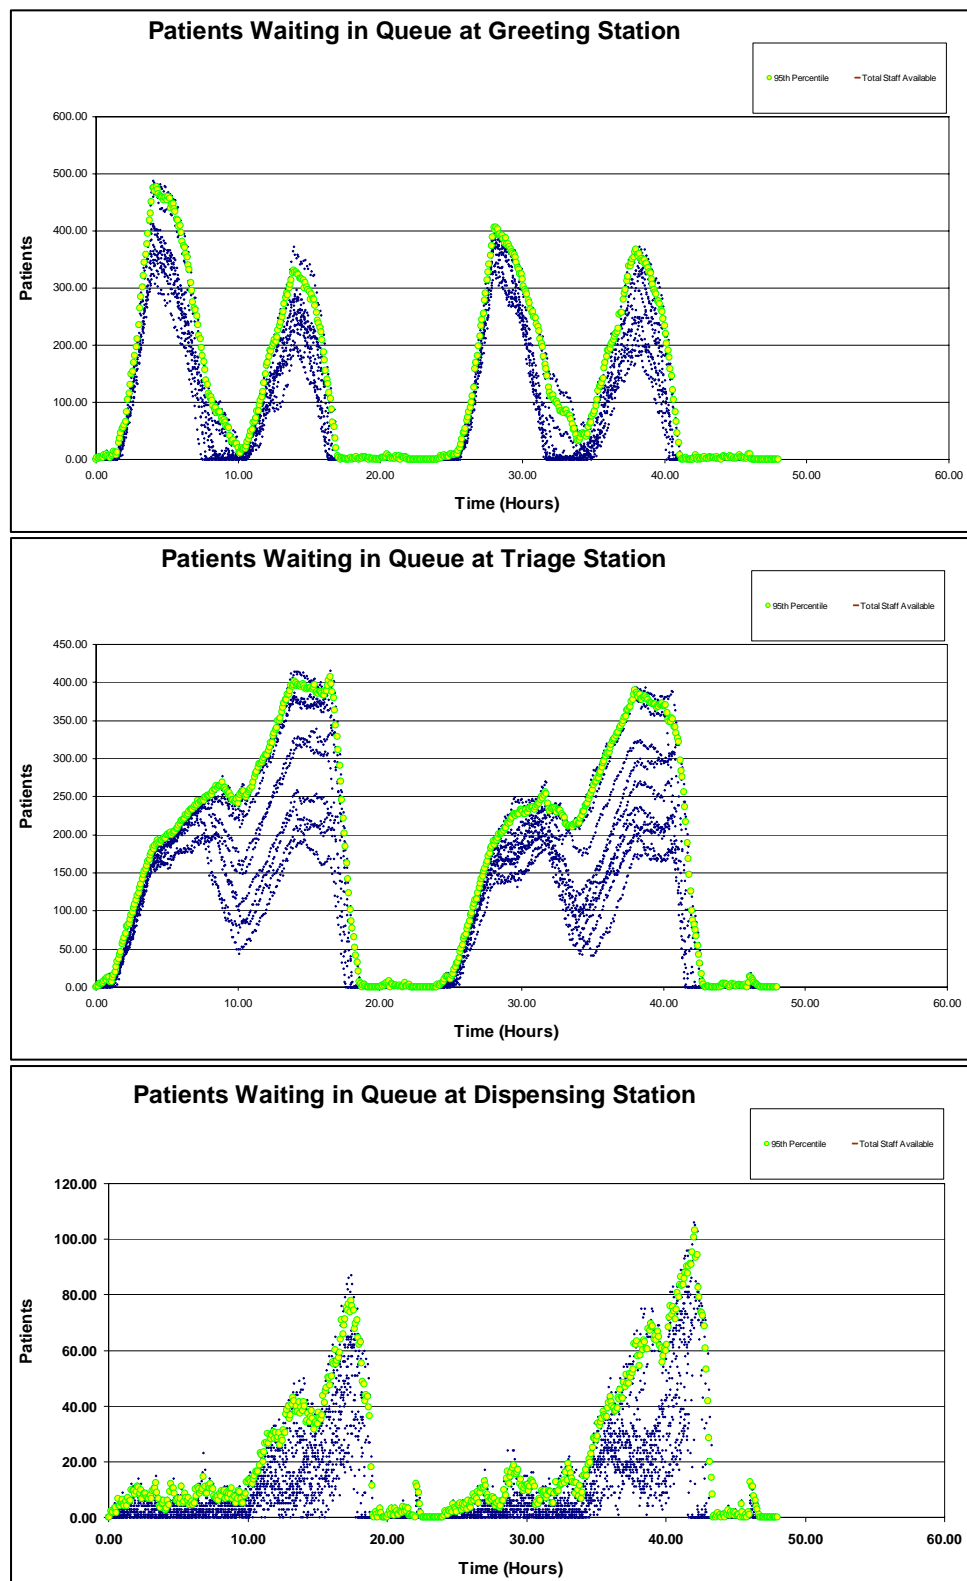


Figure 10: Queue Lengths at the POD Stations Under Patient Arrival Scenario C and Staffing Plan 3.

There are two very interesting features of the plots in Figure 10. First, notice that the queues grow to be quite large at all three stations. So, despite using only 50 fewer staff hours than staffing Plan 1, queues build up at not just one, but at all three of the stations that most patients must visit. The significant one hour delay in adjusting for staffing, and the assumption that patient arrivals will not change over the next hours, result in staffing levels that do not adequately accommodate the actual patient arrivals. Thus, we observe fast-growing queues, followed by quick drops in queue length after the staffing levels have finally been adjusted. At some points the stations appear to be stable; for instance, the first peak in the daily arrival pattern did not negatively affect the Drug Dispensing station thanks to the bottleneck that occurred at the Greeting Station. However, we can clearly see that, if a staffing plan similar to Plan 3 were implemented at a POD with a highly variable patient arrival pattern, queues would inevitably become unstable throughout the POD. Implementing staffing Plan 2 similarly results in large queues throughout the POD, although they are less severe than the queues shown in Figure 10 for Plan 3.

Note also the extremely large variance in the simulation replications. Because of uncertainty in the timing of patient arrivals and the length of service times, the queue lengths vary by hundreds of people at the Triage and Greeting stations. The consequence of this uncertainty is that even more staff must be employed than we might have predicted in order to keep patient waiting times low with a high degree of certainty. Alternatively, if the number of staff cannot be increased, then queue lengths may grow extremely large.

Thus, it is clear that the quality of planning forecasts used in setting a staffing plan could have a tremendous impact on patient waiting times and queue lengths. Of course, staffing policy is not solely affected by one's ability to predict arrival patterns or bring staff in quickly; the overall design of the POD network affects the number of staff that will be necessary to run it. We explore this issue in the next section.

### ***3. How are optimal staffing levels affected by the number of PODs in a POD network?***

We investigated the staffing consequences of having many small PODs, as opposed to fewer large PODs. Opening a large number of small-scale PODs may seem desirable; this might allow PODs to be placed in locations that are more easily accessible to patients. Also, the organization of each individual POD might be simpler. However, we will show that, if staff-hours are limited, then opening many small-scale PODs could have very negative consequences for patient waiting times and queue lengths.

We will consider a simple example. Suppose that there are 10 small PODs, to which an average of 500 patients arrive each hour. We will compare this network to one large POD which sees an average arrival rate of 5,000 patients per hour. Both systems will experience the same total expected daily patient demand. For each system, we calculated the staffing levels necessary to limit average patient waiting times to five minutes at each station within the PODs.

We find that each of the small PODs uses 984 staff-hours to provide an average patient waiting time of 6.72 minutes; this gives a total of 9,840 staff-hours to operate the POD network. The large POD, on the other hand, requires only 8,664 staff-hours to provide an average patient waiting time of 6.28 minutes. Thus, we see that almost 15 percent more staff-hours are required to operate a system of small PODs, which provides slightly worse service to patients. These staff-hour calculations do not even account for the additional overhead that would likely be required to operate a POD network that is more widely-distributed over many locations.

There are, of course, many factors that affect capacity decisions for PODs, but it is a fact that increasing the number of PODs also increases the number of staff-hours required to provide a desired level of service to the same number of patients. Clearly, then, if the number of staff-hours is limited, having a single, huge POD would provide the best service for patients. In large emergencies, such an arrangement would likely be infeasible. However, such a network could be approximated by employing an effective command and control system. If this system were well-run and infrastructure put in place that could move staff or patients between PODs, the command and control system would create a virtual single, large POD, where staffing levels and demands for service across the entire network could be monitored continuously and adjustments made as required. As we have demonstrated, the number of staff required to meet patient needs in a timely manner will be reduced substantially if such a virtual system can be created.

#### ***4. What are the consequences of using a command and control system to monitor and control the POD network?***

In the previous section we emphasized the benefits of having a command and control system in place to manage a network of PODs effectively and efficiently. Such a system would allow managers to dynamically assess staffing levels at all PODs and respond to observed patient demand by transporting staff or patients throughout the network. That is, if patient demand is especially high at one POD, more staff could be moved from a less-busy POD to help, or patients could be bused to other PODs with lower waiting times. Of course, if we can predict patient arrival rates with great accuracy and staff accordingly, then the value of such measures would be reduced. However, there is so much uncertainty regarding where and when spikes in patient demand will arise that a command and control system should be considered essential to the successful execution of a mass prophylaxis campaign.

We have described experimental results that show that patient waiting times and queue lengths are minimized in an environment in which POD staffing levels are appropriately calculated to match average patient arrival rates. These results strongly suggest that a responsive and robust POD system can exist only if mechanisms are in place that can allocate resources or patients throughout a campaign. If such a command and control environment is not created, one of two outcomes will occur. Either significantly larger staffs will have to be employed or patients will have to wait for very long times to be served. Neither of these two outcomes is desirable. We also note that these observations are not limited to the setting of staff levels. Greater amounts of other resources such as medical supplies, equipment, and space will all be needed if scarce resources are allocated ineffectively. Thus, creating a command and control system that balances loads throughout the POD network by moving patients or personnel among PODs is essential for managing system operations effectively.



## V. Conclusions

We have introduced a new POD simulation tool called D-PODS and used it to demonstrate that non-stationary and uncertain patient arrival patterns have a significant impact on POD staffing; that flexible staffing policies can greatly improve patient waiting times and queue lengths; that POD networks require fewer staff-hours when fewer large PODs are used, rather than many small PODs; and that a command and control system is essential for managing a POD network effectively. The best way to limit patient waiting times while using the fewest possible staff would be use a single huge POD and accurately predict average patient arrival rates over time. However, neither of these would likely be possible in a real emergency. Therefore, a command and control system should be used to approximate these aims, which it does by continuously monitoring all patient demand and staffing levels throughout the POD network and calculating when and where patients or staff should be moved within single PODs or between PODs within the network. Finally, it should be noted one last time that, while we have primarily discussed how D-PODS can be used to explore POD staffing policy, the observations we make can also be applied to the allocation of other resources, such as medical supplies and equipment.

## REFERENCES

1. Koh HK, Elqura LJ, Judge CM, et al. Implementing the cities readiness initiative: lessons learned from Boston. *Disaster Med Public Health Prep.* Mar 2008;2(1):40-49.
2. Wein LM, Craft DL, Kaplan EH. Emergency response to an anthrax attack. *Proc Natl Acad Sci U S A.* Apr 1 2003;100(7):4346-4351.
3. Hupert N, Mushlin AI, Callahan MA. Modeling the public health response to bioterrorism: using discrete event simulation to design antibiotic distribution centers. *Med Decis Making.* Sep-Oct 2002;22(5 Suppl):S17-25.
4. Hupert N, Cuomo J, Callahan M, Mushlin A, Morse S. *National Guidelines for mass prophylaxis using dispensing/vaccination centers (DVCs).* Washington, D.C.: DHHS Agency for Healthcare Research and Quality (AHRQ); April 2003.
5. Hupert N, Cuomo J. Bioterrorism and Epidemic Outbreak Response Model (BERM). <http://www.ahrq.gov/news/press/pr2003/btmodpr.htm>.
6. Lee E, Maheshwary S, Mason J, Glisson W. Large-Scale Dispensing for Emergency Response to Bioterrorism and Infectious-Disease Outbreak. *Interfaces.* 2006;36(6):591-607.
7. Aaby K, Abbey RL, Herrmann JW, Treadwell M, Jordan CS, Wood K. Embracing computer modeling to address pandemic influenza in the 21st century. *J Public Health Manag Pract.* Jul-Aug 2006;12(4):365-372.
8. Buzacott J, Shanthikumar J. *Stochastic Models of Manufacturing Systems.* Prentice-Hall, New Jersey. 1993.
9. Baccam P, Boechler M. Public health response to an anthrax attack: an evaluation of vaccination policy options. *Biosecur Bioterror.* Mar 2007;5(1):26-34.
10. Brookmeyer R, Johnson E, Bollinger R. Public health vaccination policies for containing an anthrax outbreak. *Nature.* Dec 16 2004;432(7019):901-904.
11. Wein LM, Craft DL. Evaluation of public health interventions for Anthrax: a report to the secretary's council on Public Health Preparedness. *Biosecurity and Bioterrorism.* 2005;3(4):348-356.
12. Wilkening DA. Sverdlovsk revisited: modeling human inhalation anthrax. *Proc Natl Acad Sci U S A.* May 16 2006;103(20):7589-7594.
13. Wilkening DA. Modeling the incubation period of inhalational anthrax. *Medical Decision Making.* Jul-Aug 2008;28(4):593-605.
14. Gross D, Harris CM. *Fundamentals of Queueing Theory.* Wiley, New York. 1998.
15. Centers for Disease Control and Prevention (CDC). 2006. *Continuation Guidance for Cooperative Agreement on Public Health Preparedness and Response for Bioterrorism--- Budget Year Six.*
16. Buzacott J, Shanthikumar J. *Stochastic Models of Manufacturing Systems.* Prentice-Hall, New Jersey. 1993.

## **Appendix A. User's Manual**

### ***1. Introduction***

The Dynamic Point of Dispensing Simulator (D-PODS) is a tool designed to help local public health officials evaluate the impact of policy on the operations of a single Point of Dispensing (POD) during a mass prophylaxis campaign. D-PODS is a Monte-Carlo simulation model that accepts user inputs describing a POD's layout, patient arrival patterns, and staffing plan, and then outputs a probabilistic assessment of system performance over time. The goal of this user manual is to explain the features and limitations of D-PODS as well as to offer advice on using D-PODS to successfully model POD systems.

D-PODS is implemented in Visual Basic and its user interface is an Excel workbook. A large number of user inputs are required to run the simulation. These are entered by the user in a sequence of four worksheets. Once the simulation is run, a number of output statistics and plots are automatically generated by D-PODS. This guide will explain how to get started with D-PODS and then how to enter information on the user input sheets. The guide will further identify potential pitfalls and assumptions implicit in the model that may cause surprising results. Finally, the guide will discuss how the simulation outputs can be analyzed and interpreted, and how D-PODS can be used to explore policy decisions.

## 2. Glossary of Important Terms

1. *Arrival Rate* – the hourly rate at which patients arrive to the POD (this can vary throughout each day).
2. *Arrival Types* – different groups of patients who may require different service times based on their special needs. These might include a single person, a family, an elderly person, a non-English speaker, etc.
3. *Maximum Average Waiting Time* – the maximum patient waiting time in minutes, assuming that patients arrive at the expected rate. This is an input that is used by the D-PODS staffing calculator to estimate reasonable staffing levels.
4. *Queue Length* – the total number of people waiting to be served at a given location.
5. *Routing Probability* – the probability that a patient will be routed from one station to the next station in the layout of the model. The routing probabilities are sometimes called “transition rates.”
6. *Service Time Distribution* – the probability distribution that governs the length of time that will be required to serve a single patient at a work station. These distributions are assumed to be triangular and hence are defined by minimum, most likely (mode), and maximum values.
7. *Service Time Increase Factor* – the percent increase in service time required by different arrival type groups. The service times for different arrival types are calculated by drawing a random service time from the appropriate service time distribution and then multiplying this number by one plus the service time increase factor. For example, if the service time increase factor is 100%, the service time doubles for that arrival type. The factor is the same for all service stations presently.
8. *Simulation Replication* – one repetition of the simulated prophylaxis campaign.
9. *Staff Utilization* – the average percentage of staff serving customers. This value is weighted by time.
10. *Time Period* – the duration of time over which the expected patient arrival rate to the POD is constant, also referred to as an “arrival interval.”
11. *Worker Interval* – the duration of time for which the number of staff is constant at all stations.

### 3. Working with D-PODS

To use D-PODS to model a particular POD, a larger number of user inputs will be required. These data are entered in a series of worksheets. At the beginning of each section below we provide a checklist that identifies what you will need to complete the relevant steps of the model.

#### 3.1 Getting Started

*You will need:*

1. *Excel 2003 or later*
2. *The files “PoD\_DB.mdb” and “PoD\_v27.xls”*

Make sure that you have both the files “PoD\_DB.mdb” and “PoD\_v27.xls” stored in the same directory folder; otherwise D-PODS will not run. No other files should be placed in the D-PODS folder because they may cause errors to arise while running the program. If you ever try to run the simulation and get an error that stops the simulation from completing, then go to the directory where these folders are stored and delete all files except these two.

To get started, open the Excel file “PoD\_v27.xls”. A “Security Warning” message box may prompt you to set appropriate security settings; if this happens, select the button to “Enable Macros”. Once the file is open, ensure that the following References are selected:

1. Visual Basic For Applications
2. Microsoft Excel 11.0 Object Library
3. Microsoft DAO 3.6 Object Library
4. Microsoft Forms 2.0 Object Library
5. OLE Automation
6. Microsoft Office 11.0 Object Library
7. Microsoft ActiveX Data Objects 2.8 Library
8. Microsoft OLE DB Error Library

To do this, go to *Tools → Macros → Visual Basic Editor*. From the Visual Basic Editor window, go to *Tools → References*. Look up the above references and mark the corresponding checkboxes. Then click “Okay” and return to the Excel – POD\_v27 window.

### 3.2 Starting the Model

The file opens to a cover page. Click the “Begin Model” button in the lower right-hand side of the screen, and then the “Case Selection” page will appear, as shown in Figure 1. The purpose of this page is to allow you to load the inputs from a previous case or to create a new case.



Figure 1: Case Selection Screen Shot

If you select the “Create New” radio button, a message box will pop up to ask if you would like to continue. If you answer “Yes”, you are taken to the “D-PODS Menu” page. If you answer “No”, then the program simply remains on the “Case Selection” page and awaits your next decision.

If you select the “Existing Case” radio button, a user form will pop up, showing a list of all cases currently in the database.

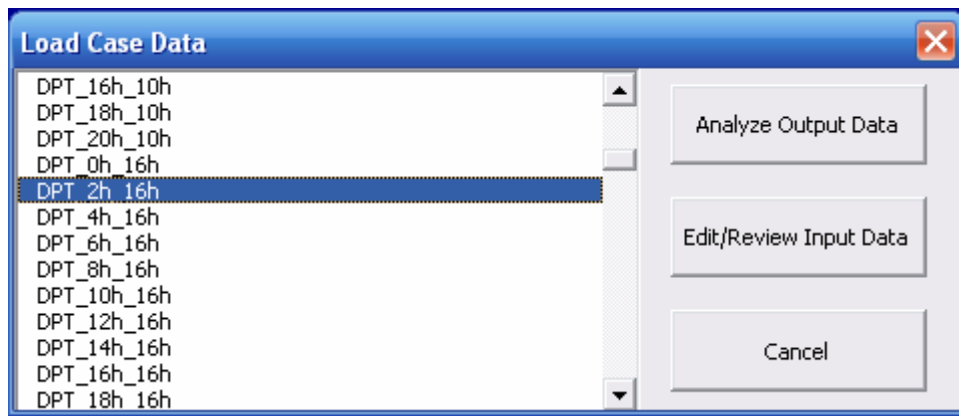


Figure 2: Load Case Data Screen Shot

After selecting the case of interest, you can choose to “Analyze Output Data” or “Edit/Review Input Data”. You can also click “Cancel”, which returns the screen to the “Case Selection” page. Once you have chosen your case, you can click on either button. Both buttons will load the inputs and simulation outputs of the case that you chose. The only difference is that clicking “Edit/Review Input Data” will bring you to the “D-PODS Menu” page, while clicking “Analyze Output Data” button will bring you directly to the “Output Tables” page.

### 3.3 D-PODS Menu

*You will need:*

1. To fill out the information in steps 5 and 6. This will be discussed in greater detail in section 3.8.

The “Main Menu” shows all of the different choices you have when you are working with D-PODS and suggests the best order in which to perform each task. To complete steps 1 through 4, you must click on their respective buttons and fill in the information on the worksheet that appears. Steps 5 through 7 are concerned with the details of saving inputs and running the simulation. Step 8 takes you to various sheets that display the simulation output.

## D-PODS Menu

**In order to run the program, follow the steps as shown below.  
The program may not work if executed in the incorrect order.**

<b>Step 1</b>	Construct the model	Model
<b>Step 2</b>	Input arrival rate	Arrivals
<b>Step 3</b>	Input the service time parameters	Service
<b>Step 4</b>	Establish staffing levels	Staffing
<b>Step 5</b>	Enter simulation parameters	
	Number of simulation replications	2
	Random Seed	150
<b>Step 6</b>	Input Case Name	[ Enter Name ]
<b>Step 7</b>	Run the simulation	Run
<b>Step 8</b>	View the Results. The results are displayed in both tabular and graph form.	
		Output Tables
		Output Graphs

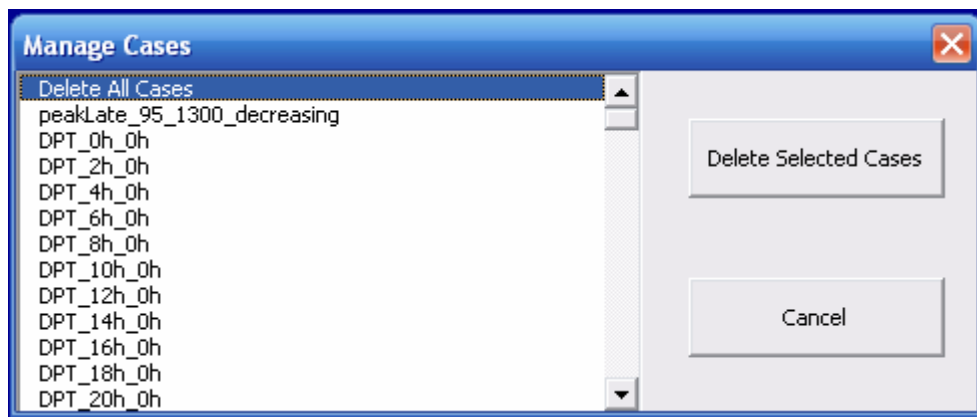
**Other Options**

Select an existing case.	Existing Case
Manage case list.	Manage Cases
Go to Cover Page to Start Over	Start Over

**Figure 3: D-PODS Menu Screen Shot**

The buttons at the bottom of the menu allow you to work with the database file or to restart the program. If you click the “Existing Case” button, you will be prompted with the same user form as on the “Case Selection” page. You can then select a different case to study.

If you click the “Manage Cases” button, you will be prompted by a user form very similar to the one shown when “Existing Case” is clicked. The difference is that this button allows you to delete cases. This allows you to manage the size of the database and to clean out old data that are no longer of interest.



**Figure 4: ManageCases Screen Shot**

Before a case is deleted from the database, a warning box will appear to ask whether you are sure that you want to delete the case. If you select “Yes”, the case is deleted. If you select “No”, nothing happens and you will be returned to the “D-PODS Menu” page.

The “Start Over” button will bring you back to the cover page and will allow you to start the program from scratch.



### 3.4. Step 1: Constructing the Model

You will need to know:

1. How many days you would like to simulate
2. How many hours the POD will operate each day (not including overtime)
3. How many workstations will be in the POD and the names of these stations
4. The station transition probabilities. That is, given the current location of a patient, the probability that he will next move to each other station
5. The total capacity of the POD

When you click the “Model” button on the D-PODS menu, the “Construct the Model” worksheet will appear, which is illustrated in Figure 5.

**Step 1: Construct the Model**

**Step A**

Input duration of the campaign in days	2
Input hours of operation per day	22

**Step B**

Number of Stations (maximum is 15)	4
Input station names and exits	Station Names
Input station transition probabilities	Transition Rates

**Step C**

Input PoD capacity	50000
Return to Main Page	Return

**Station Names Table**

Station	1	2	3	4	
Name	Greeting	Triage	Medical Eval	Drug Dispensing	Exit

**Transition Rates Table**

	Greeting	Triage	Medical Eval	Drug Dispensing	Exit
Greeting		0.95	0.05	0.00	0.00
Triage			0.05	0.95	0.00
Medical Eval				0.95	0.05
Drug Dispens					1.00

The first station is the immediate station after entering the PoD. All subsequent stations should be listed in a downward flowing manner. Meaning, if station j can be reached from station i, then i should be listed before station j. Likewise, one must exit from the last station.

The transition rate matrix below indicates the probability of flowing from state i to state j within the PoD. The Exit column in the transition rate matrix below the probability of flowing out of the PoD from any of the stations.

Figure 5: Step 1 Screen Shot

To complete this input step, complete Steps A, B, and C in order. In Step A, you first specify the duration of the simulation in days. You may want to begin by considering only one or two days, rather than the entire duration of the planned prophylaxis campaign, because the simulation run time is proportional to the number days simulated; doubling the number of days approximately doubles the simulation run time. Also, it is important to note that D-PODS assumes that all days are identical. That is, the POD is required to have the same operating plan and patient arrival pattern for each day. Beyond random variation, the only reason why two days might have significantly different performances is if patient arrivals from one day cannot be served by the end of the day and must be helped during the following work day. We discuss this further below.

The “hours of operation per day” is the period of time during which patients are allowed to arrive and enter the POD. In Figure 5, this value is set to 22. That means that any patients who arrive to the POD between the 22<sup>nd</sup> and 24<sup>th</sup> hour of the day will be turned away without service. However, any patients who have already entered the POD will be served as long as staff is available, even if this means that the POD remains open for longer than 22 hours. It is possible that not all of these patients will be served by the time patients begin arriving again for the next day; in this case, the newly arrived patients simply queue behind the waiting patients. On the last day of the simulation, however, the simulation will stop at the correct time, even if some patients have not been served.

In Step B you describe the layout of the POD by declaring how many workstations will be in use and their names. D-PODS assumes that each station has one queue where patients wait for service. There can be 0 or more staff working at each station, and the person at the front of the queue will be served by the next available staff person at that station. When the patient has finished receiving service, he immediately enters service at the next station or the queue for that station, if no server is currently available. Travel time between stations is assumed to be insignificant.

To complete Step B, first enter a whole number for the “number of stations” and then click the “Station Names” button. This generates the “Station Names Table.” Be careful, however, because clicking this button also clears current entries from the table, and the “Undo” function in Excel will not let you get the information back. If you want to save the information, copy it to another location before clicking the button. You can type the names of each station. It is important to enter the station names in the correct order, because it is not possible for patients to travel from higher-numbered stations to lower-numbered stations. So, if you name the first station “Greeting” and the second station “Triage,” then it will be possible (but not required) for patients to move from Greeting to Triage, but it is not possible to go from Triage to Greeting. Note that this makes it impossible for patients to visit any station in the POD more than once. So, the first station listed in the “Station Names Table” should correspond to the immediate station that patients encounter upon entering the PoD. The last station in the POD is a “dummy” station called Exit, which indicates where patients leave the POD after receiving treatment.

After completing the “Station Names Table”, you can click on the “Transition Rates” button; this will generate the “Transition Rates Table”, which allows you to input patient routing probabilities through the POD. So, for example, the entry in row 3, column 4 of the table gives

the probability that a patient at station 3 will move to station 4 after she has finished receiving service at station 3. Since all patients must move to a new station after service, the probabilities in each row must sum to one. If the values do not add up to one, a warning box will appear when you try to move to a new worksheet. Notice that some cells in the “Transition Rates Table” are grayed out; this enforces the condition described above, making it impossible to move from a higher-numbered station to a lower-numbered one. It also makes it impossible for a patient to remain at the same station.

In Step C, you can enter the POD capacity, which sets the maximum number of patients allowed in the POD at any point in time. The capacity may depend on the size of the building or it may be set low to limit the risk of transmitting disease. When the number of patients in the building exceeds the POD capacity, patients form a line outside of the POD and wait until another patient leaves the building before entering the queue for the first patient inside the POD. Since these patients may be turned away without being counted at the end of the day and they are not included in statistics describing the average time in the POD, some of the output statistics may be misrepresentative if the queue outside the POD grows large. In the experiments described in the main paper, the POD capacity was set to be almost five times larger than the expected number of daily patient arrivals, so that all patients who arrived to the POD immediately entered the queue for the first station.

Once Steps A – C are complete, you can proceed to “Step 2: Input Arrival Information.” There are two ways to do this. One option is to return to the “D-PODS Menu” page by clicking the “Return” button in Step D and then clicking the “Arrivals” button on the menu page. The other option is to simply click the “Generate Arrivals” button on the left side of the screen. The font on the button of the “Construct the Model” step is grayed out to help you remember where you currently are in the process.

### **3.5. Step 2: Input Arrival Rate Information**

Step 2 allows you to enter information describing patient arrivals throughout each day. Figure 6 shows the “Input Arrival Rate Information” worksheet.

In Step A, you first state how many patient arrival intervals you want to use to describe the daily patient arrival pattern. You will be able to set a new expected patient arrival rate for each arrival interval, but during each interval the expected patient arrival rate remains constant. The number

of arrival rate intervals per day, can be anywhere between 1 and 96. Choosing 96 as the value means that each arrival interval will be 15 minutes long, because there are 24 hours in each day.

Step 2: Input Arrival Rate Information						
<b>Step A</b>	Input number of arrival intervals	12				
<b>Step B</b>	Input start time (Must be input in time format)	7:00 AM				
<b>Step C</b>	Input number of arrival types. Arrival types include single person, families, elderly, non english speakers, etc.	2				
<b>Step D</b>	Generate chart based on information provided in steps A and B	Chart				
<b>Step E</b>	Fill in the highlighted cells of the chart on the right based on the headers					
<b>Step F</b>	Return to Main Page	Return				

Interval Number	Hours per Interval	Interval Start	Interval End	Arrivals Per Hour: Type 1	Arrivals Per Hour: Type 2
1	2.00	7:00 AM	9:00 AM	300	150
2	2.00	9:00 AM	11:00 AM	200	100
3	2.00	11:00 AM	1:00 PM	500	250
4	2.00	1:00 PM	3:00 PM	500	250
5	2.00	3:00 PM	5:00 PM	700	350
6	2.00	5:00 PM	7:00 PM	400	200
7	2.00	7:00 PM	9:00 PM	200	100
8	2.00	9:00 PM	11:00 PM	100	50
9	2.00	11:00 PM	1:00 AM	100	50
10	2.00	1:00 AM	3:00 AM	50	25
11	2.00	3:00 AM	5:00 AM	0	0
12	2.00	5:00 AM	7:00 AM	200	100

Figure 6: Step 2 Screen Shot

In Step B, you can input the start time, which is the time the POD opens each day. If you want to simulate a POD scenario in which patients begin arriving before the POD opens, you can set the staffing levels at the beginning of the day to 0 at each station in Step 4: Establish Staffing Levels. If you want to simulate a POD that opens before patients begin coming, enter “0” for the patient arrival rates at the beginning of the day in Step E.

Step C allows you to enter the number of patient arrival types. Your basic arrival type may be single, mobile, English-speaking adults, but other useful arrival types may include disabled individuals, school children, non-English speakers, or any other groups of special interest who require service times different from the average. In Step 3, it will be possible to define different service rates for the different patient arrival types.

Step D requires that you click the “Chart” button to generate the patient arrival rates table. Then, in Step E you can define the arrival intervals and the patient arrival rates during each time interval. The first light blue column requires the duration of each arrival interval in hours; these durations do not all need to be the same, but each one must be rounded to the nearest 0.25 hour and the values must sum to 24 hours. The next two columns show the start and end time of the interval. The subsequent columns allow you to enter the expected patient arrival rates during each interval. Any blank spots in the chart will be treated as 0’s. Patients arrive to the POD according to a nonhomogeneous Poisson process, which means that the interarrival times

between subsequent patients are distributed according to a Poisson distribution. The mean of the Poisson process at each point in time is calculated based on the expected values that you enter in this chart.

Note that the arrival rates can only be entered for a single day, because we assume that the arrival pattern on the subsequent days is identical. As mentioned earlier, if the expected arrival pattern varies by day, then a simulation must be run for each individual day.

After you have completed the table, you may proceed to Step 3: Service Time Parameters. To do this, either click “Return” to go back to the “D-PODS Menu” page and then click the “Service” button, or simply use the “Set Service Parameters” button on the left side panel.

### 3.6. Step 3: Service Time Parameters

Step 3 allows you to input parameters that govern how long it takes to serve each patient type at each station. Figure 7 below shows a screen shot of the Step 3 worksheet.

Step 3: Service Time Parameters																														
<b>Step A</b>	Generate chart to input percentage increases in service time and chart to input parameters of service time for a single person for each station.	<input type="button" value="Chart"/>	<table border="1"> <thead> <tr> <th>Arrival Type</th> <th>Service Time Increase Factor</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>0%</td> </tr> <tr> <td>2</td> <td>20%</td> </tr> </tbody> </table>	Arrival Type	Service Time Increase Factor	1	0%	2	20%	<table border="1"> <thead> <tr> <th>Station</th> <th>Minimum</th> <th>Most Likely</th> <th>Maximum</th> </tr> </thead> <tbody> <tr> <td>Greeting</td> <td>0.25</td> <td>0.50</td> <td>1.00</td> </tr> <tr> <td>Triage</td> <td>1.00</td> <td>2.00</td> <td>3.00</td> </tr> <tr> <td>Medical Eval</td> <td>2.50</td> <td>5.00</td> <td>10.00</td> </tr> <tr> <td>Drug Dispensing</td> <td>0.50</td> <td>1.00</td> <td>2.00</td> </tr> </tbody> </table>	Station	Minimum	Most Likely	Maximum	Greeting	0.25	0.50	1.00	Triage	1.00	2.00	3.00	Medical Eval	2.50	5.00	10.00	Drug Dispensing	0.50	1.00	2.00
Arrival Type	Service Time Increase Factor																													
1	0%																													
2	20%																													
Station	Minimum	Most Likely	Maximum																											
Greeting	0.25	0.50	1.00																											
Triage	1.00	2.00	3.00																											
Medical Eval	2.50	5.00	10.00																											
Drug Dispensing	0.50	1.00	2.00																											
<b>Step B</b>	Fill in the cells of the chart on the right with the percent increase. The first arrival type (single person) is the baseline and has zero percent increase in service time.																													
<b>Step C</b>	Fill in the cells of the chart on the far right with the parameters of service time for a single person (in minutes).																													
<b>Step D</b>	Return to Main Page	<input type="button" value="Return"/>																												

Figure 7: Step 3 Screen Shot

The service time at a given station for some patient is a random number drawn from that station’s service time distribution, which is assumed to be triangular. D-PODS further assumes that each patient is processed individually and that the service times for each station and arrival type remain constant throughout the simulation; there is no accounting for increases or decreases in speed due to worker fatigue or learning curve.

In Step A, you must click the “Chart” button to generate the two charts on the right side of the screen. The structure of this chart is based on the stations you defined in Step 1 and on the number of arrival types you entered in Step 2. Default charts may be displayed before you click

this button, but it is important that you click the “Chart” button to ensure that the tables display the correct numbers of arrival types and stations.

Step B requires that you fill in the Service Time Increase Factor table. Choose one of your arrival types to be the base arrival type; enter “0” or “0%” for this arrival type’s service time increase factor. The service time increase factors for the other arrival types describe how the processing time of the other arrival types compare to your base arrival type. For example, if Arrival Type 1 is your base arrival type, with a service time increase factor of 0 and the service time increase factor for Arrival Type 2 is 20%, then the service times for Type 2 patients will be 20% longer than the service times for Type 1 patients, on average. If Type 2 patients have the same service time distributions as Type 1 patients, then you should set the Type 2 service time increase factor to 0%. If Type 2 patients are processed 20% faster than Type 1 patients, then for the Type 2 service time increase factor you should enter “-20%”. All numbers in this table must be between -100 and 100. Note, however, that D-PODS does not currently allow you to define an arrival type that is served faster than Type 1 at some stations and slower than Type 1 at other stations; the service time increase factor is applied to all of the stations.

To complete Step C, you must fill out the far right table. Enter the parameters that describe the triangular service time distributions at each station for Arrival Type 1. The minimum and maximum times are strict limits on the service times, and the most likely time is the peak of the triangular distribution, which means it is the mode of the service times. Note that the minimum time must be less than or equal to the most likely time, and the most likely time must be less than or equal to the maximum time. If all three values are the same, then the service time will be constant at that station.

After completing both tables, you may proceed to Step 3: Service Time Parameters. To do this, either click “Return” to go back to the “D-PODS Menu” page and then click the “Staffing” button, or simply use the “Establish Staffing Levels” button on the left side panel.

### 3.7. Step 4: Staffing Requirements

The fourth input step is setting the staffing plan at each station throughout each day. Figure 8 shows a screenshot of the Staffing Requirements sheet.

Step 4: Staffing Requirements													
<b>Step A</b>	Input number of worker intervals per day		<b>8</b>										
	Start time from arrivals page		<b>7:00 AM</b>										
<b>Step B</b>	Generate staff interval chart		<b>Chart</b>										
<b>Step C</b>	Input number of hours per worker interval												
<b>Step D</b>	Calculate the maximum arrival rate per hour for each worker interval		<b>Arrival</b>										
<b>Step E</b>	Fill in the desired maximum average queue waiting time in minutes in the chart below												
<b>Step F</b>	Calculate advised number of workers per interval based on the max in Step D		<b>Advice</b>										
<b>Step G</b>	Fill in the highlighted columns for the desired number of staff at each location for each worker interval. If you'd like to use the advice data, click the button to the right.		<b>Use advice?</b>										
<b>Step H</b>	Return to Main Page		<b>Return</b>										

Worker Interval	Hours per Interval	Interval Start Time	Interval End Time	Maximum Arrival Rate	Max Avg Waiting Time	Greeting Advice	Greeting Input	Triage Advice	Triage Input	Medical Eval Advice	Medical Eval Input	Drug Dispensing Advice	Drug Dispensing Input
1	3	7:00 AM	10:00 AM	450	5	5	5	16	16	9	9	10	10
2	3	10:00 AM	1:00 PM	750	5	8	8	26	26	13	13	16	16
3	3	1:00 PM	4:00 PM	1050	5	11	11	36	36	17	17	22	22
4	3	4:00 PM	7:00 PM	1050	5	11	11	36	36	17	17	22	22
5	3	7:00 PM	10:00 PM	300	5	4	4	11	11	7	7	7	7
6	3	10:00 PM	1:00 AM	150	5	2	2	6	6	5	5	4	4
7	3	1:00 AM	4:00 AM	75	5	1	1	3	3	4	4	2	2
8	3	4:00 AM	7:00 AM	300	5	4	4	11	11	7	7	7	7

Figure 8: Step 4 Screen Shot

Step A allows you to set the maximum number of times during the day that you want to change the number of staff at each station. If you plan to keep the same number of staff at each station all day long, you may enter a 1 here. The period of time during which staffing levels at each station are constant is called a “worker interval.” D-PODS does not allow worker intervals shorter than two hours, so the maximum number of worker intervals each day is 12. That is, you may not change the number of staff at each station more than 12 times each day.

The second line shows you the start time for the POD operations, which you entered in Step 2: Arrival Rate Information.

Step B involves clicking the “Chart” button to generate a new chart at the bottom of the page. However, note that clicking this button will clear the information currently in the table, and this

cannot be undone. Be sure to save any important information about your previous work before clicking “Chart.” If the number of worker intervals you entered in Step A is the same as the number of intervals currently shown in the chart, you do *not* need to click the “Chart” button.

Step C requires that you enter the duration of each worker interval in the “Hours per Interval” column of the table. This column must total 24 hours and each interval must be at least two hours long, but the intervals do not have to be of equal lengths. If the intervals do not sum to 24 hours, then a message will appear at the bottom of the table. Notice that the next two columns of the table display the start and end times of each interval.

To complete Step D you must click the “Arrival” button to display the maximum expected hourly patient arrival rate during each interval in the “Maximum Arrival” column. This information is taken directly from the patient arrivals table in Step 2: Input Arrival Information. Even if you do not plan to use this information in developing your staffing plan, you must click the “Arrival” button, or a warning will appear that prevents you from proceeding to the next step.

In Step E, you must fill in the “Max Avg Waiting Time” column of the table. The numbers in these columns indicate the maximum expected waiting time for each station that you would like to have during each segment of the day. It is not possible to set different maximum expected waiting times for different stations. Once you have entered these values, which must be positive real numbers, you must click the “Advice” button in Step F. The D-PODS staffing calculator then uses a calculation from Buzacott<sup>1</sup> to estimate the number of staff required at each station to limit the maximum average patient waiting times. These estimates are shown in “Advice” columns of the table. Note that D-PODS *requires* that you complete Steps E and F, even if you do not plan to use this information in your staffing plan.

Step G requires you to fill in the “Input” columns of the table with the number of staff at each station during each worker interval. One way that you may choose to do this is by using the “Use Advice?” button, which copies the values from the “Advice” columns into the “Input” columns. If you click the “Use Advice?” button, a warning will pop up to ask if you are sure that you want to use the advice. If you click “Yes” to indicate that you really do want to use the advice from the D-PODS staffing calculator, any numbers currently in the “Input” columns will be replaced by the advice numbers, and it is not possible to undo this choice.

Once all of the “Input” columns are complete, you should return to the “D-PODS Menu” page by clicking “Return” or by clicking the “Menu” page on the left side panel. It is also possible to



immediately begin running the simulation by clicking the “Run” button on the left side panel, but if you do this you will not have a chance to save your inputs or modify the simulation parameters.

### **3.8. Step 5: Simulation Parameters**

To complete Step 5, you must input the simulation parameters on the “D-PODS Menu” page. The number of simulation replications is the number of complete prophylaxis campaigns you would like to simulate. If, for example, you set your campaign to be 3 days long and you declare that you want to run 100 replications, D-PODS will simulate the 3 day campaign 100 times. Note that doubling the number of simulation replications will double the total run time. If you are not sure how fast D-PODS will run on your computer, start off with only a few simulations so you don’t crash Excel. If your computer is fairly new and fast, you should have no trouble doing 10 to 20 or more runs of D-PODS at one time. The number of simulation replications should be a positive integer; if you enter a non-integer number, it will automatically be rounded.

The “Random Seed” also must be entered in Step 5. This number should be a positive integer. It is used in the Visual Basic random number generator to determine the random numbers used in the simulation. If you use the same random seed with exactly the same set of inputs, you will get the same output. However, if you use the same set of inputs but change the random seed, the simulation output will be different (although the average values will likely be quite similar). If you want to compare different inputs and keep the comparison as “fair” as possible, then it is a good idea to use the same random seed for both simulations. This helps eliminate random error as a confounding factor in your comparison results.

Step 5 can be completed at any point during the input process, but it is not optional. The simulation will not run if either light blue box is empty.

### **3.9. Step 6: Input Case Name**

In Step 6 you have an opportunity to save the inputs that you have entered and the simulation output that will be generated when you run the simulation. In the light blue box that initially reads “[ Enter Name ]” click and type the name you would like to use. If you do this, you will be able to reopen this case using the Select Existing Case dialog. If you forget to enter a case name before running the simulation and you want to save your work, you have two options. One is to return to the D-PODS menu, enter a case name in the Step 6 box, and then re-run your

simulation. Another option is to save the entire Excel workbook under a new name, but if you do this you may not be able to save all of the simulation output.

### **3.10. Step 7: Run the Simulation**


To execute the simulation, you must click the “Run” button in Step 7. The simulation run time depends primarily on the number of simulations runs you have chosen; the number of days in your prophylaxis campaign, and the number of patient arrivals. Large simulations may easily take up to 5 or 10 minutes to complete. If more than 10 minutes go by, or if you think that the run time is too long, you can stop the simulation by hitting the “Esc” key. Then click “End” in the error dialog box that pops up. If you do this, you *must* go to the directory where you have saved D-PODS and delete any new files that have been generated; there will likely be one called “Simulation\_Output.txt”. If you do not delete this file, D-PODS will not run in the future. You may then modify your simulation by decreasing one or more of the parameters discussed above to speed up the running time.

You will know that the simulation is complete when the mouse pointer changes from an hourglass back into an arrow.

### 3.11. Step 8: View the Results

D-PODS produces a variety of output tables and graphs that you can use to analyze POD performance. From the D-PODS menu, click on either the “Output Tables” button or “Output Graphs” button to begin viewing the results of your simulation. You can easily move between the two using the buttons on the left side panel or by returning to the D-PODS menu at any time.

Clicking the “Output Tables” button will bring you to the “Output Tables” worksheet, as shown below in Figure 9.



Cornell University

# Output Tables

Construct Model	Number of Arrivals	6,028	6,028	0.00
	Total Throughput	6,028	6,028	0.00
	Average Throughput (pt./min.)	2.09	2.09	0.00
Generate Arrivals	Average Time in POD (min./pt)	5.37	5.37	0.00
	Average Number in POD (pt./min.)	11.24	11.24	0.00
Set Service Parameters	Average Time in System (min./pt.)	5.37	5.37	0.00
	Average Number in System (pt./min.)	11.24	11.24	0.00
Establish Staffing Levels	<b>Performance Measures at Each Station by Worker Interval:</b> <div> Select Worker Interval: <div></div> </div>			

**Figure 9: Output Tables Initial Screen Shot**

The tables at the top of the sheet show summary statistics that describe the entire prophylaxis campaign. The number of patient arrivals is shown, as is the total patient throughput. Patient throughput is the number of patients who were processed during the campaign. If all patients who arrived were processed before the end of the campaign, these two numbers will be the same. The means and standard deviations are taken across all of the simulation replications.

“Average Throughput” is calculated for each simulation replication by dividing the total throughput by the operating hours. The mean and standard deviation of these numbers are then reported. “Average Time in POD” is calculated for each simulation replication by taking the mean of the patient times in the POD. The numbers reported are calculated from these averages. “Average Number in POD” is the average number of patients in the POD at a given point in

time. The values reported in the output table are the sample mean and standard deviation of these numbers.

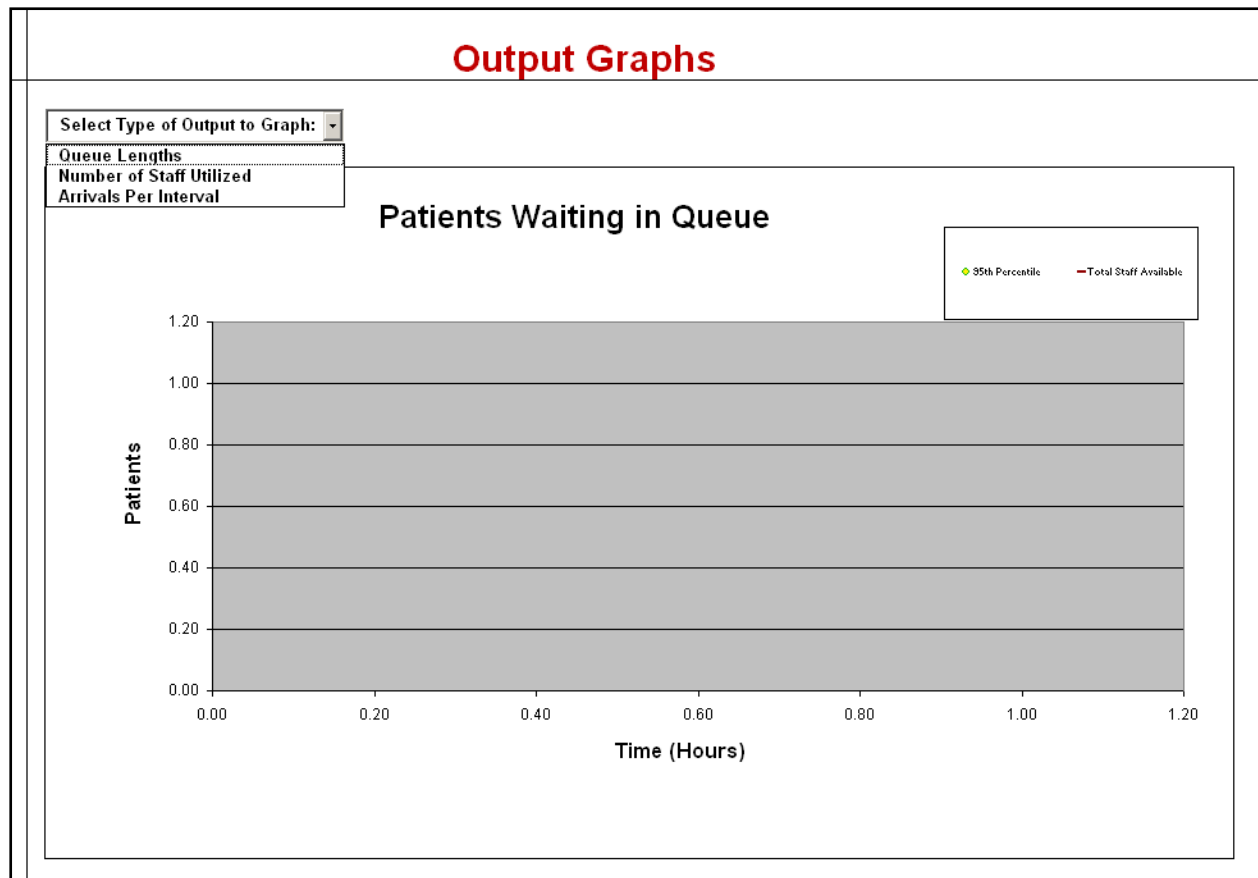
“Average Time in System” and “Average Number in System” are similar to the previous two statistics, except that they include the patients waiting outside the POD, so these means will always be at least as large as the “... in POD” statistics. If there are never patients waiting outside the POD, these values will be equal to “Average Time in POD” and “Average Number in POD,” respectively.

The statistics are further broken down by workstation and by worker intervals. You can use the drop down box beneath the main output table to select either the entire campaign or a particular worker interval. For the time period chosen, you can see statistics that describe the queue lengths, patient queue times, and staff utilization at each station.

Performance Measures at Each Station by Worker Interval:			
4			
	Estimate*	Mean	Std Dev
<b>Queue Outside of POD:</b>			
Average Queue Length (patients)	0.00	0.00	0.00
Maximum Queue Length (patients)	0.00	0.00	0.00
Average Wait in Queue (minutes)	0.00	0.00	0.00
<b>Station 1</b>			
Average Queue Length (patients)	0.08	0.14	0.07
Maximum Queue Length (patients)	1.00	1.50	0.71
Average Wait in Queue (minutes)	0.07	0.07	0.00
Maximum Staff Used	2.00	2.00	0.00
Staff Utilization	0.45	0.51	0.08
<b>Station 2</b>			
Average Queue Length (patients)	0.63	0.85	0.32
Maximum Queue Length (patients)	5.00	6.50	2.12
Average Wait in Queue (minutes)	0.36	0.47	0.15
Maximum Staff Used	4.00	4.00	0.00
Staff Utilization	0.78	0.78	0.00
<b>Station 3</b>			
Average Queue Length (patients)	0.00	0.00	0.00
Maximum Queue Length (patients)	0.00	0.00	0.00
Average Wait in Queue (minutes)	0.00	0.00	0.00
Maximum Staff Used	3.00	3.50	0.71
Staff Utilization	0.18	0.22	0.05
<b>Station 4</b>			
Average Queue Length (patients)	0.19	0.20	0.01
Maximum Queue Length (patients)	2.00	2.00	0.00
Average Wait in Queue (minutes)	0.07	0.08	0.01
Maximum Staff Used	3.00	3.00	0.00
Staff Utilization	0.54	0.59	0.07

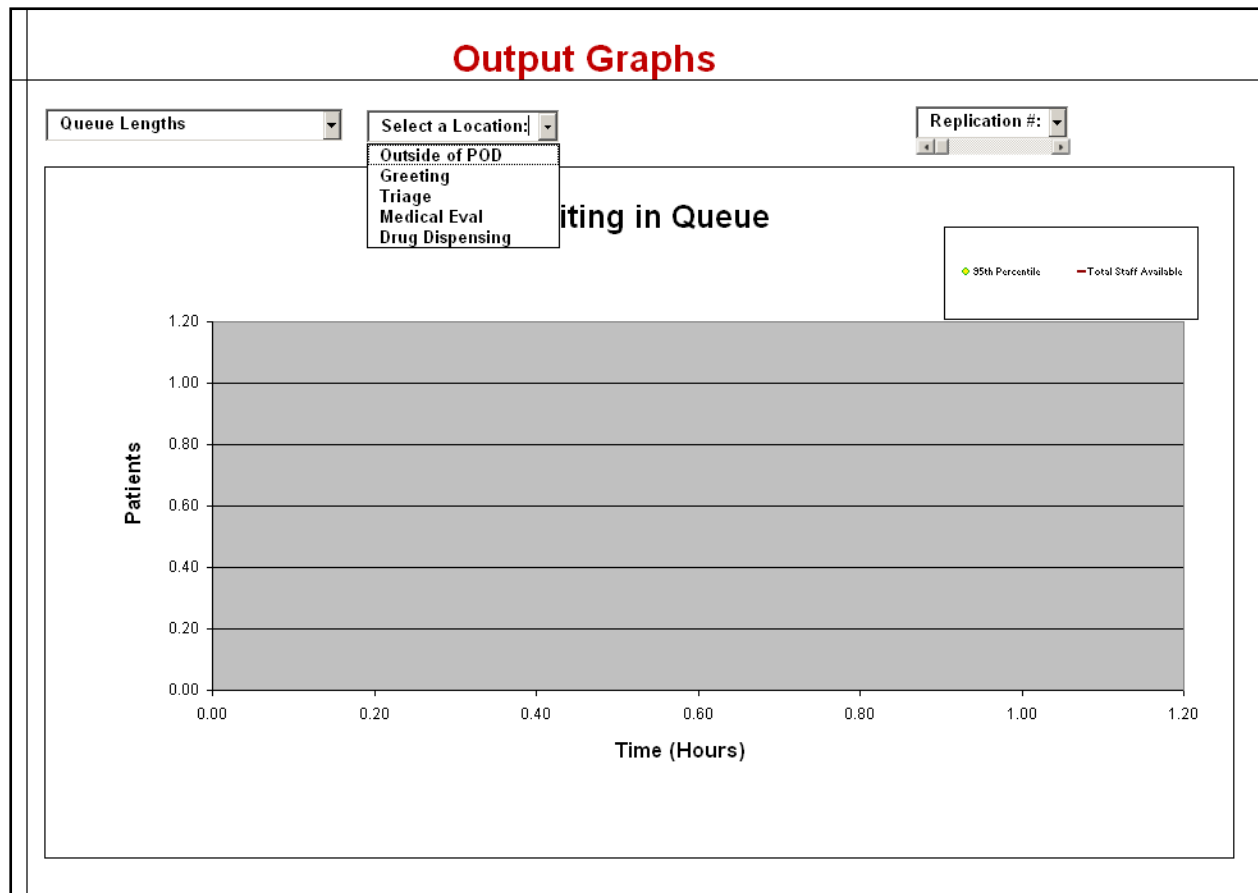
**Figure 10: Output Tables Station Performance Measures Screen Shot**

Clicking the “Output Graphs” button on either the left panel from the “Output Tables” page or on the “D-PODS Menu” page brings you to the “Output Graphs” page. A screen shot of this page is shown below in Figure 11.



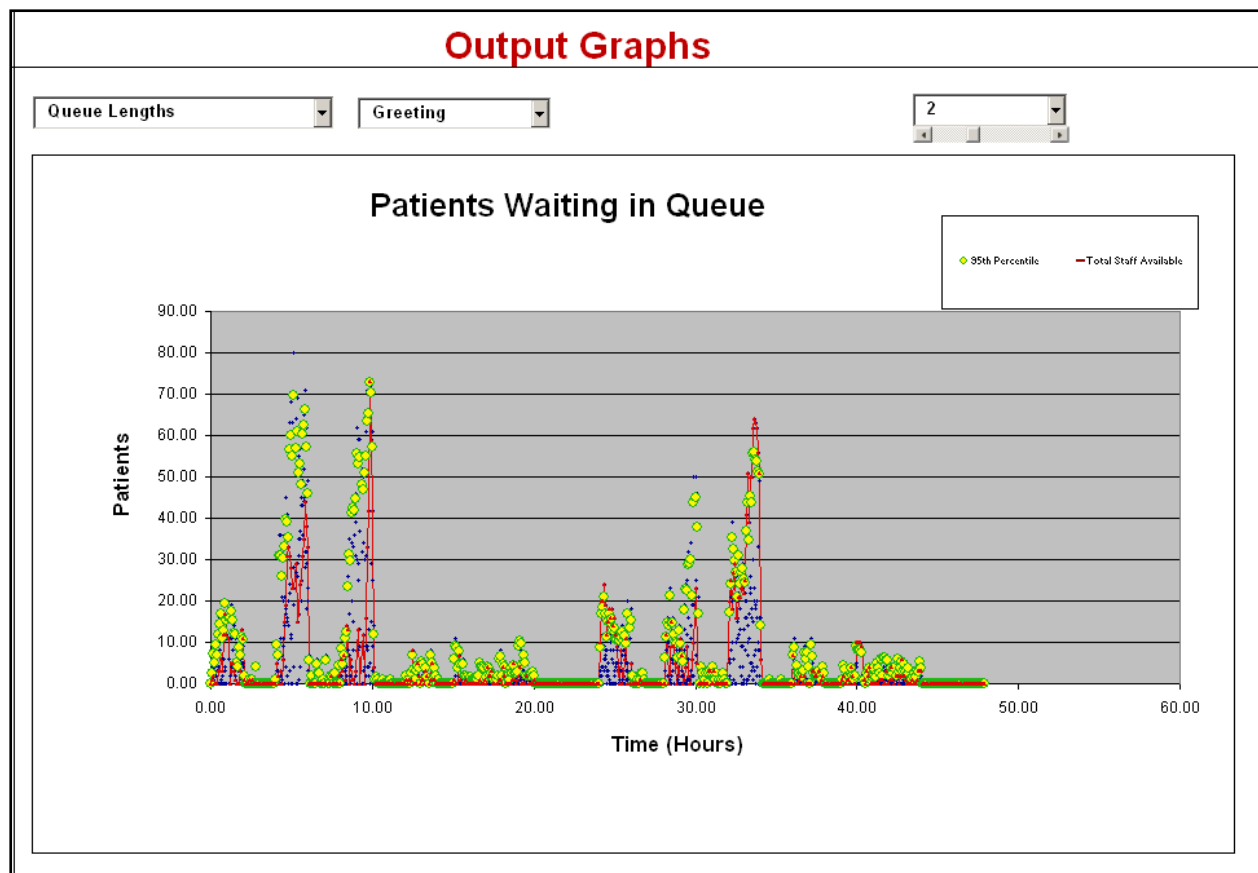
**Figure 11: Output Graphs First Selection Screen Shot**

Only one graph can be displayed at a time. The drop down box (expanded in Figure 11) allows you to choose the type of data you would display: “Queue Lengths,” “Number of Staff Utilized,” or “Arrivals Per Interval.” Once you have made a selection, two new boxes will appear. The first allows you to select the location for which you would like to see information. As shown in the expanded drop down box in Figure 12, this can be any of the work stations in the POD or “Outside of POD.” However, note that if your POD is large enough to accommodate all arriving patients, then there will not be any non-zero data for the “Outside of POD” option.



**Figure 12: Screen Shot of Output Graphs Select a Location**

Once you have entered valid options in each of the first two boxes, a plot of the data you have selected will be shown, as in Figure 13.



**Figure 13: Screenshot of Output Graphs after Location Selection and Replication Selection**

Next, you may select a specific simulation replication from the drop down box on the right, or you may choose to view all replications. The scroll bar below the drop down box also controls the simulation replication to be selected. After selecting a specific simulation replication, the data for that replication are highlighted with a red line; in Figure 13, replication 2 is shown.

The data shown in these plots are drawn from the Access database. In order to plot the data in different ways or view the raw numbers, you must explore the database directly. This is discussed in the next section.

#### ***4. Navigating the Access Database***

When you click the “Run” button to start your simulation, D-PODS immediately saves all input data to the database. The input data are stored in the tables titled “CaseList”, “SingleInputs”, “ArrivalRates”, “ServiceIncreaseFactor”, and “Staffing”.

The “CaseList” table stores the case name specified on the “D-PODS Menu” page. This table also assigns the unique CaseListID to each CaseName, which creates a relationship throughout the tables. All data are stored with the CaseListID included as one of the Access fields. This makes creating queries relating to a specific case name much easier.

The “SingleInputs” table contains all input values that are not entered through a chart within D-PODS, such as the number of simulation replications or the number of days in the prophylaxis campaign. It also contains the station names and station transition probabilities. The values are stored based on the CaseListID, the variable name (varName), and the value the variable holds in the program (varValue).

The “ArrivalRates” table contains the chart shown on the right on the “Step 2” page in D-PODS. The data are organized in the Access table just as it is in D-PODS. The “ServiceIncreaseFactor” table stores the information from the chart on the “Step 3” page in D-PODS that details the service time increase factor for each of the arrival types. The “Staffing” table details the information held in the chart at the bottom of the “Step 4” page in D-PODS.

While the simulation is running, the output data are stored in a text file called “Simulation\_Output.txt”, which is transferred to the database at the end of the simulation. These data are moved into the Access tables titled “SingleOutputs” and “Outputs”. The “SingleOutputs” table is organized exactly like the “SingleInputs” table and contains information that is displayed in the tables on the “Output Tables” page in D-PODS, such as the number of arrivals for the duration of the planned campaign. The “Outputs” table contains the output data that are stored at the end of every simulated five-minute interval. These data include the replication number, the shift number, the event time, the size of the queues at each location, the number of arrivals to each location within the interval, the average wait time for each location, and the number of servers that are busy for each location. These data are organized by CaseListID and can easily be queried in order to do further analysis on the output data.



## 5. Designing POD Staffing Policy Experiments

In this section we explain exactly how to recreate the staffing levels used in our experiments. We took advantage of the D-PODS staffing calculator, which provides staffing level estimations are based on queuing approximations found in Buzacott.<sup>1</sup> Our approximations output staffing levels given two inputs: an expected patient arrival pattern and an average patient waiting time for each station,  $W$ . If the patient arrival rate during the staffing interval is constant in expectation, then D-PODS will calculate staffing levels such that the average patient waiting time will approach  $W$  in the limit as the number of patients served increases to infinity.

Of course, patient arrival rates may not be constant over an entire staffing interval. In this case, the D-PODS staffing calculator conservatively estimates the arrival rate using the maximum patient arrival rate during the interval. This has the effect of making the average patient waiting time lower than the maximum allowed,  $W$ .

For staffing plan 1, we used the D-PODS staffing calculator described above to determine the workforce requirements for each station over time. For staffing plan 2, we imagined that one half hour before each staffing change, we could observe the true average patient arrival rate to the POD,  $L$ . We then used the same Buzacott queuing formula mentioned above to determine the staffing levels that would be required to set the average patient waiting time to five minutes at each station, assuming that patients arrived at a constant rate of  $L$  patients per hour. We set the workforce requirements for the next two hour staffing interval to equal these levels.

To better understand staffing plan 2, consider the following example. Suppose that at 4:00 pm, we can change the number of staff working at each station in the POD; this may involve calling in new staff or sending people home, as well as shifting workers between stations. However, we need to declare our decisions to the workforce by 3:30, so that current workers can prepare to move around or go home and new staff have time to travel to the POD. So, just before 3:30, we determine the average current patient arrival rate for the past half hour; suppose that on average, we have been seeing 50 patients arrive every five minutes, which is equivalent to 600 patients per hour. We then use the Buzacott queuing formula to determine the optimal staffing levels for each station in the POD, assuming that patients continue to arrive at an average rate of 600 patients per hour. At 6:00 pm we will have a chance to reassign and rearrange staff once again, so we will repeat this process just before 5:30 pm.

Staffing plan 3 was similar to plan 2, except that we set staffing levels based on the patient arrival rates that occurred one hour before each staffing change. So, in the example above, if we were set to have a staffing change at 4:00 pm, we would calculate staffing levels just before 3:00 pm. So, if an average 55 patients were arriving every five minutes from 2:30 to

3:00 pm, we would calculate new staffing levels assuming that patients continue to arrive at an average rate of 660 patients per hour (which is equivalent to 55 patients per five minutes).

In staffing plan 4, we suppose that we could perfectly forecast the patient arrival rates at the beginning and the end of the staffing interval. We took the average of these two numbers, and set the workforce levels at each station according to the Buzacott formula, assuming that patients arrived at a constant rate equal to this average for the entire staffing interval. In the example above, this would mean that, before the staffing change must be made at 4:00 pm, we could look at our forecasting mechanism and see exactly the average arrival rate from 4:00 pm to 4:30 pm and the average arrival rate from 5:30 pm to 6:00 pm. We would then average these two numbers and use the average to calculate staffing levels using the Buzacott formula.

## **6. References**

1. Buzacott J, Shanthikumar J. *Stochastic Models of Manufacturing Systems*. Prentice-Hall, New Jersey. 1993.

## Appendix B. Future Event Lists

A typical Monte Carlo Simulation uses a single future event list (FEL) to keep track of events that are scheduled to occur, such as arrivals and departures from different queues and service locations. The events in each FEL include only those which are necessary to complete the next step in the simulation. For example, since the next arrival time only depends on the current arrival time (based on a distribution of inter-arrival times), the only event stored in the FEL is the current arrival; all of the previous arrival times can be deleted from the FEL. The next arrival is not scheduled (and placed into the list) until the current arrival is executed (and removed from the list).

Future event lists are sorted as new events are scheduled and must be inserted into the list. The earliest event is executed, and removed from the top, or head, of the list.

In order to decrease the run time for our model, we chose to organize our simulation in a different way. Instead of having a single FEL, we created one FEL for each service station, plus one more to deal with arrivals outside of the POD. Thus, if there are  $n$  stations, there are  $n+1$  FELs. The size of the FEL for a given station at some point in time is therefore limited to the current number of servers at that station. At each point in time, the number of departures is limited by the number of servers. This is because once the station has reached capacity, one cannot know the time of the next departure until one person completes service. With  $n+1$  FELs, each list remains quite small and can be searched quickly when new items must be inserted. Of course, operating multiple FELs in a time sequence requires adding an extra step to determine which FEL has the earliest event at its head. Of course, this simply means that the heads of the FELs must be compared, which requires little time. The simulation then processes that event. The simulation continues to operate as a typical Monte Carlo simulation by performing the tasks relating to the event found.

Thus, changing a single FEL to  $n+1$  smaller FELs reduces the time for inserting an event, with a slight compromise in additional time for removing an event.